



Barracuda Link Balancer



Copyright Notice

Copyright 2004-2011, Barracuda Networks

www.barracuda.com

v2.2-110503-01-0503

All rights reserved. Use of this product and this manual is subject to license. Information in this document is subject to change without notice.

Trademarks

Barracuda Link Balancer is a trademark of Barracuda Networks. All other brand and product names mentioned in this document are registered trademarks or trademarks of their respective holders.

Contents

Chapter 1 – Introduction 1

Overview	1
About this Guide	2
Features of the Barracuda Link Balancer.	2
Link Management	3
Aggregating Link Bandwidth	3
Link Failover.	3
Outbound Link Load Balancing	3
Inbound Link Balancing and Failover	3
VLAN Support	4
High Availability	4
Persistence	4
Bandwidth Management and Quality of Service (QoS)	4
Traditional Firewall	4
Site-to-Site VPN and Link Failover	5
Ability to Deploy with Your Network Firewall	5
Local Network Services	6
Reporting	6
Web User Interface	6
Technical Support.	6

Chapter 2 – Installing the Barracuda Link Balancer 7

Deployment Scenarios	7
In Front of Firewall	8
Replacing Your Firewall	9
Overview of the Installation Steps	10
In Front of Your Firewall Installation	11
Replacing Your Firewall Installation	15

Chapter 3 – Configuring the Barracuda Link Balancer 19

Configuring Network Settings.	19
Adding, Updating or Viewing WAN Link Configuration	19
Adding a New WAN Link	19
Adding Static Routes	20
Configuring VLANs	20
Creating IP Aliases	21
Configuring DNS Servers	21
Configuring Per Interface Health Checks	21
Configuring the DHCP Server.	21
Configuring the Firewall	21
Firewall Functionality	22
Order of Execution of Firewall Rules	22
Inbound Firewall Rules	22
Inbound 1:1 NAT Rules	22

Port Forwarding Rules	23
Outbound Firewall Rules	23
Firewall Logging.	23
Creating Custom Applications	24
Managing Bandwidth	24
Link Usage for Inbound and Outbound Traffic	24
Creating Bandwidth or Quality of Service (QoS) Rules	25
Outbound Traffic Routing	25
Specifying the Link Used by Outgoing Traffic	26
Changing the Source IP Address of Outgoing Traffic	27
Configuring Virtual Private Networks	27
Site-to-Site VPN Tunnels	27
Creating VPN Tunnels	28
Creating a VPN in a NAT'd Environment	28
Failover and Failback	29
VPN Tunnel as Failover Link for a Broken Site-to-Site WAN Link	29
Troubleshooting a VPN Tunnel	29
Configuring the DNS Server for Inbound Load Balancing	30
Introduction	30
DNS Records Time to Live	30
Recommended Deployment	31
Split DNS	31
DNS Zone Transfer Blocking	31
Becoming an Authoritative DNS Host	31
If You Add a WAN Link After the Domains are Created	34
Zones and Domains.	34
DNS Records	35
Configuring Administrative Settings	35
Controlling Access to the Web User Interface	36
Changing the Default Password	36
Setting Email Addresses for Alerts	36
Customizing the Appearance of the Web User Interface	36
Setting the Time Zone of the System	36
Enabling SSL for Administration	37

Chapter 4 – Creating a High Availability Environment 39

Overview	39
Ethernet Passthrough	39
Operation of High Availability (HA)	39
Physical Connectivity of the Clustered Systems	40
Requirements for Clustered Systems	40
Synchronization of Data Between Clustered Systems	41
Failover and Failback	41
Planning Your High Availability Deployment	42
In Front of Single Network Firewall	42
In Front of Dual Network Firewalls	43
No External Firewalls	44
Creating a Cluster	44
Removing a System from a Cluster	46
Updating Firmware on Clustered Systems	46

Chapter 5 – Monitoring the System 47

Checking Status	47
Viewing Logs	47
Using a Syslog Server to Centrally Monitor System Logs	48
SNMP Monitoring	48
SNMP Traps	49
System Reports	49
Viewing System Tasks	50

Chapter 6 – Maintaining the Barracuda Link Balancer 51

Backing up and Restoring Your System Configuration	51
Updating the Firmware of Your Barracuda Link Balancer	51
Replacing a Failed System	51
Reloading, Restarting, and Shutting Down the System	52
Using the Reset Button to Reset the LAN IP address	52
Using the Built-in Troubleshooting Tools	53
Rebooting the System in Recovery Mode	53
Reboot Options	54
Barracuda Networks Limited Hardware Warranty (v 2.1)	55
Exclusive Remedy	55
Exclusions and Restrictions	55
Barracuda Networks Software License Agreement (v 2.1)	56
Barracuda Networks Energize Updates and Other Subscription Terms	60

Chapter 1

Introduction

This chapter provides an overview of the Barracuda Link Balancer and includes the following topics:

<i>Overview</i>	1
<i>Features of the Barracuda Link Balancer</i>	2
<i>Technical Support</i>	6

Overview

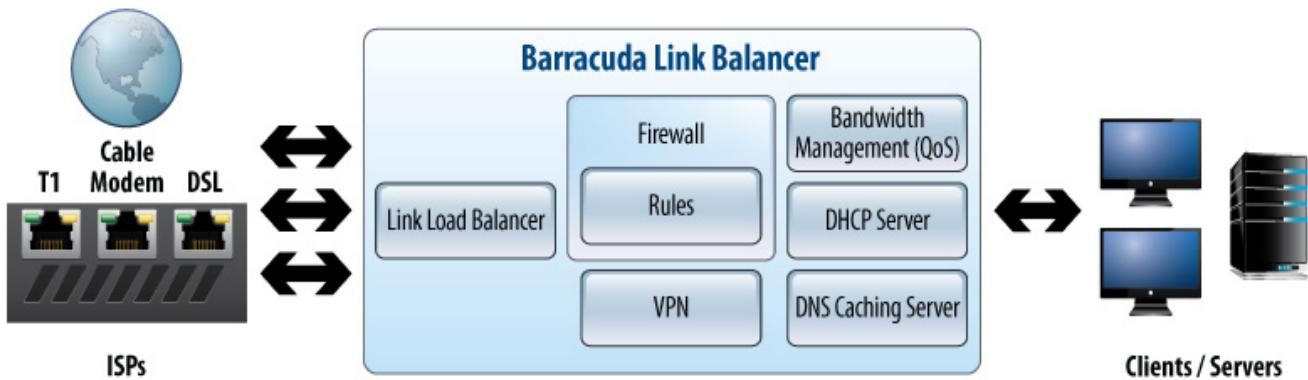
The Barracuda Link Balancer routes and manages traffic across multiple Internet connections or WAN links. By using multiple inexpensive connections from one or more Internet service providers, you can reduce the need to purchase high speed and high cost links. Supported links include T1, T3, E1, DSL, cable, fiber optic and MPLS. The Barracuda Link Balancer:

- balances incoming and outgoing network traffic across multiple links
- provides automated failover in case of link failure
- manages bandwidth
- performs Quality of Service (QoS) for Internet applications
- includes DHCP server and DNS caching server functionality
- can act as a traditional firewall or can be installed in front of your existing firewall
- provides site-to-site VPNs with link failover and fallback.

The Barracuda Link Balancer is *not* designed for load balancing that distributes incoming traffic across servers. If you are interested in that functionality, the Barracuda Load Balancer may meet your needs.

As shown in [*Figure 1.1*](#), the Barracuda Link Balancer provides an interface between multiple Internet connections and your clients and servers.

Figure 1.1: Barracuda Link Balancer Functionality



About this Guide

This guide provides a general discussion of the features and concepts that you need to understand in order to best configure the Barracuda Link Balancer. Other than the installation instructions that are in *Installing the Barracuda Link Balancer* on page 7, you will find that detailed, procedural configuration steps are in the online help of the Web user interface.

When referring to specific feature settings, this guide specifies the name of the tab in the Web user interface followed by a (>) and the actual page name. For example, you can view link utilization and performance statistics on the **Basic > Status** page.

Features of the Barracuda Link Balancer

This section describes the features of the Barracuda Link Balancer:

<i>Link Management</i>	3
<i>Aggregating Link Bandwidth</i>	3
<i>Link Failover</i>	3
<i>Outbound Link Load Balancing</i>	3
<i>Inbound Link Balancing and Failover</i>	3
<i>VLAN Support</i>	4
<i>High Availability</i>	4
<i>Persistence</i>	4
<i>Bandwidth Management and Quality of Service (QoS)</i>	4
<i>Traditional Firewall</i>	4
<i>Site-to-Site VPN and Link Failover</i>	5
<i>Ability to Deploy with Your Network Firewall</i>	5
<i>Local Network Services</i>	6
<i>Reporting</i>	6
<i>Web User Interface</i>	6

Link Management

The Barracuda Link Balancer can manage links that have static or dynamic (DHCP) IP addresses and can authenticate using PPPoE.

Aggregating Link Bandwidth

The Barracuda Link Balancer automatically aggregates Internet bandwidth from multiple links to the same or diverse sources. Administrators can choose multiple links to the same or different ISPs for the purposes of consolidating access to affordable Internet bandwidth.

Any single session (e.g. a TCP stream) has at most only the bandwidth from any one WAN link. One computer may have more than one session if it is connected to more than one remote site.

Link Failover

The Barracuda Link Balancer regularly checks the health of each Internet link and only uses the available links. If it detects a link failure, the failed link is removed from link balancing. When the failed link becomes available again, the Barracuda Link Balancer will resume using that link. All of this happens without administrator intervention.

If a link fails, existing sessions on that link will be disconnected. Clients who were using the failed link will be able to reconnect quickly to their destination using another available link rather than having to wait for the original link to be restored.

Outbound Link Load Balancing

When traffic from a client IP address going to a new destination IP address is detected, the Barracuda Link Balancer selects which link to use. It calculates the available capacity for each link based on uplink speed and current usage and uses the link with the largest available capacity. If needed, you can create outbound routing rules to override this behavior.

Inbound Link Balancing and Failover

The Barracuda Link Balancer uses authoritative DNS to direct incoming connections to a WAN link. When an external user accesses a Web site, for example, that is hosted behind the Barracuda Link Balancer, a DNS request is sent to the Barracuda Link Balancer for the IP address of the site. The Barracuda Link Balancer returns the IP address of the site which directs the traffic to a WAN link.

When determining which IP address to return, the available capacity for each link based on configured speed and current usage is calculated. The link with the largest available capacity is returned so that adaptive inbound load balancing is achieved. Also, if a link is found to have failed, the address for that link is not returned until it becomes available again.

In order to accomplish this, the Barracuda Link Balancer acts as an authoritative DNS server for the domains or sub-domains that you host. You can create DNS records on the Barracuda Link Balancer to identify your domain and to map that domain to multiple externally accessible IP addresses.

VLAN Support

The Barracuda Link Balancer supports Layer 2 VLANs.

High Availability

The Barracuda Link Balancer supports High Availability configurations where two Barracuda Link Balancers are deployed as an active-passive pair.

Persistence

The Barracuda Link Balancer automatically tracks the IP addresses of each client / source and corresponding server / destination. As long as the source and destination IP address pair are the same, traffic between them will use the same link. In addition, any one source and destination IP address pair will be tied to a specific link through about 15 minutes of inactivity. If traffic from an already tracked source IP address is detected, it may be sent on a different link if the destination IP address is unique.

Bandwidth Management and Quality of Service (QoS)

The Barracuda Link Balancer includes software that can automatically prioritize critical Internet applications. For example, you can assign priority to Web browsing and email while giving peer-to-peer applications and media streaming a lower priority. In this way, you can ensure that bandwidth-intensive applications do not interfere with business-critical operations.

Traditional Firewall

The Barracuda Link Balancer incorporates standard firewall functionality, including:

- Network Address Translation (NAT):
 - IP masquerading - Clients in the internal network are protected from the Internet. All Internet services appear to be provided by the Barracuda Link Balancer firewall, while the internal clients remain invisible.
 - 1:1 NAT - You can directly assign external addresses to internal servers. Ideal for hosting internal applications or services requiring regular outbound requests such as SMTP, 1:1 NAT provides a secure method to match additional external addresses with a single internal server for inbound and outbound traffic.
 - Port forwarding (or Port Address Translation) - The traffic to the same port across one or more multiple links is directed to an internal client.
 - Many to 1 NAT - One internal server may receive traffic from more than one WAN link. You can achieve this by creating 1:1 NAT rules or port forwarding rules.
- IP access lists - Use IP access lists to allow or deny access, either inbound or outbound, to remote networks, clients, applications, services and ports.
- Port blocking.
- Assistance in preventing and mitigating distributed denial of service attacks (DDoS).

Site-to-Site VPN and Link Failover

You can create a site-to-site VPN tunnel between two Barracuda Link Balancers or between a Barracuda Link Balancer and another device that supports IPsec. Networks connected via a tunnel will communicate as if they are on the same network, even though they are separated by the Internet.

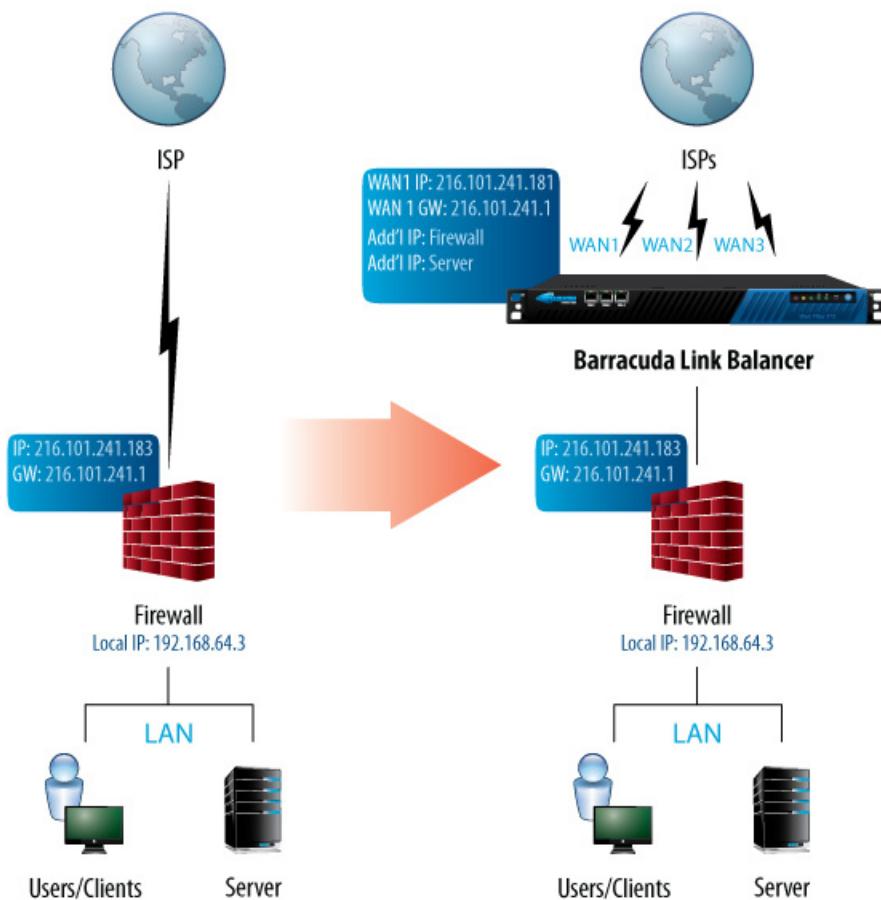
Using this functionality allows your site-to-site VPN tunnel to automatically failover to a secondary link in the case of the failure of a primary link.

Ability to Deploy with Your Network Firewall

If you already have a firewall that meets your needs, you can use the link balancing, failover and bandwidth management capabilities of the Barracuda Link Balancer and disable its firewall functionality.

As you can see in *Figure 1.2*, adding the Barracuda Link Balancer to your network without removing your firewall can be done with minimal disruption to your existing network.

Figure 1.2: Deploying with an Existing Network Firewall



Local Network Services

The Barracuda Link Balancer includes the following local network services:

- DHCP server - The Barracuda Link Balancer can automatically provision client IP addresses using the DHCP protocol. Along with defining traditional DHCP options, administrators may view active leases in real time.
- DNS caching server - The Barracuda Link Balancer caches responses to DNS queries so that repetitive DNS requests are served quickly and locally.

Reporting

A variety of trend and activity reports for the WAN links, VPNs and other system components can be generated on-demand or scheduled. Reporting is only available on models 330 and above.

Web User Interface

The Barracuda Link Balancer configuration can be administered through an SSL-secured Web user interface. Access can be through the LAN or, if configured, any of the WAN interfaces. The Web user interface can also be used for viewing traffic statistics, monitoring the health of network components and doing troubleshooting.

Technical Support

To contact Barracuda Networks Technical Support:

- By phone: call 1-408-342-5400, or if you are in the United States, (888) Anti-Spam, or (888) 268-4772
- By email: use support@barracuda.com
- Online: visit <http://www.barracuda.com/support> and click on the **Support Case Creation** link.

There is also a Barracuda Networks Support Forum available where users can post and answer other users' questions. Register and log in at <http://forum.barracuda.com>.

Chapter 2

Installing the Barracuda Link Balancer

This chapter provides instructions for installing the Barracuda Link Balancer. It includes the following topics:

<i>Deployment Scenarios</i>	7
<i>In Front of Your Firewall Installation</i>	11
<i>Replacing Your Firewall Installation</i>	15

Once you have chosen your deployment scenario, go to the section that describes how to install your Barracuda Link Balancer in that configuration. There are many steps and some are similar between the methods, but for ease of configuration they are in different sections.

Deployment Scenarios

The two most typical deployment methods are:

- *In Front of Firewall* on page 8 - Keep your existing firewall, and insert the Barracuda Link Balancer in between your firewall and the Internet. The Barracuda Link Balancer firewall is disabled in this case.
- *Replacing Your Firewall Installation* on page 15 - Replace your firewall with the Barracuda Link Balancer.

One important factor in deciding which deployment mode to choose is whether you want to replace your firewall or keep your existing firewall.

The Barracuda Link Balancer firewall provides full firewall functionality. If, however, you already have a firewall that meets your needs, you can disable the Barracuda Link Balancer firewall while still making use of the Barracuda Link Balancer's link balancing, failover and bandwidth management capabilities.

Table 2.1 shows factors to consider when choosing a deployment mode.

Table 2.1: Deployment Modes

	In Front of Firewall	Replacing Your Firewall
Network Location	The Barracuda Link Balancer is deployed between your existing firewall and the Internet.	The Barracuda Link Balancer acts as your firewall.
Barracuda Link Balancer LAN IP address	Used only for management. Can be any internal or public address that is reachable through your existing firewall from the LAN.	The default gateway for your network.
Firewall Rules	No changes to your existing firewall.	You will need to recreate any existing firewall rules on the Barracuda Link Balancer.

Table 2.1: Deployment Modes

	In Front of Firewall	Replacing Your Firewall
WAN Link	If you are enabling inbound access to resources behind the Barracuda Link Balancer, such as a Web server, at least one WAN link must have a static IP address.	The Barracuda Link Balancer may use the same IP address that had been used by your firewall.
Site to Site VPN	If you already have a site to site VPN it should be terminated on your existing firewall. VPN traffic has one source IP address so it goes out on only one WAN link. It is recognized as VPN traffic so it will not be NAT'd by the Barracuda Link Balancer. No failover or fallback is available. Alternatively, make the Barracuda Link Balancer a VPN endpoint to achieve failover and fallback to and from a secondary link.	Failover and fallback to and from a secondary link.

In Front of Firewall

Figure 2.1 gives an example of a customer network that has both client and server traffic.

Figure 2.1: Deployment Example - before Barracuda Link Balancer

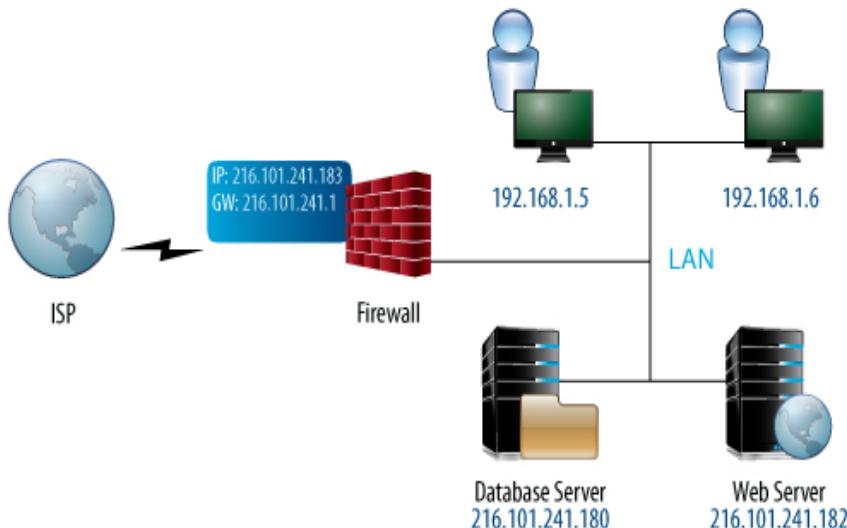


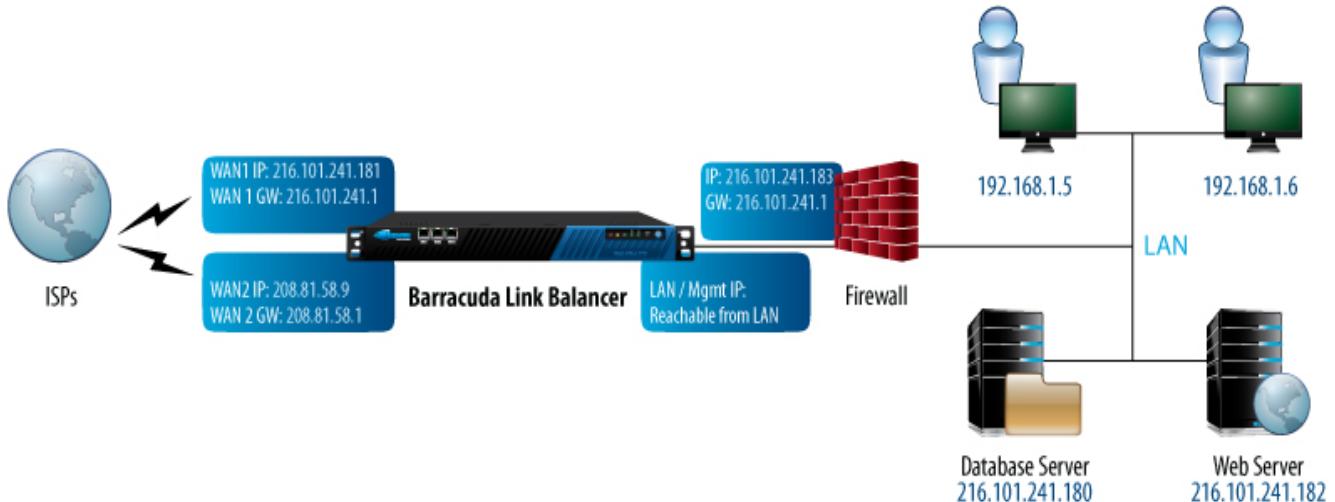
Figure 2.2 shows the same network with a Barracuda Link Balancer that was installed with no changes to the configuration of the existing firewall. A new WAN link has been added.

In this network:

- The Barracuda Link Balancer has a static IP address on WAN1 that is on the same network as the firewall and the externally visible servers.
- The clients are on a different subnet than all WAN links.
- The external IP address and gateway of the firewall remain the same.
- The gateway IP addresses of the Barracuda Link Balancer and the firewall are provided by the ISPs. The gateway of the LAN devices is provided by the firewall.

- The Barracuda Link Balancer LAN IP address can be any internal or public address that is reachable through your existing firewall from the LAN. You may allocate an external IP address for it, or choose a non-routable IP address. If the latter, it should be on a different subnet than the LAN devices already on the network. Remember that if the firewall does not recognize an address as being on the local network it will pass it to the Barracuda Link Balancer.

Figure 2.2: Deployment Example - Barracuda Link Balancer in front of firewall



Replacing Your Firewall

Figure 2.3 gives another example of a customer network that has both client and server traffic.

Figure 2.3: Deployment Example - before Barracuda Link Balancer

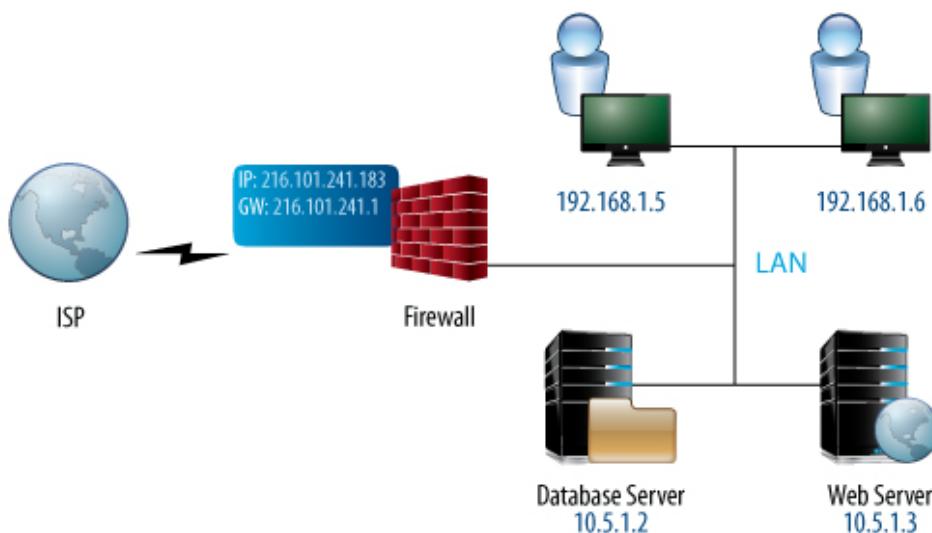


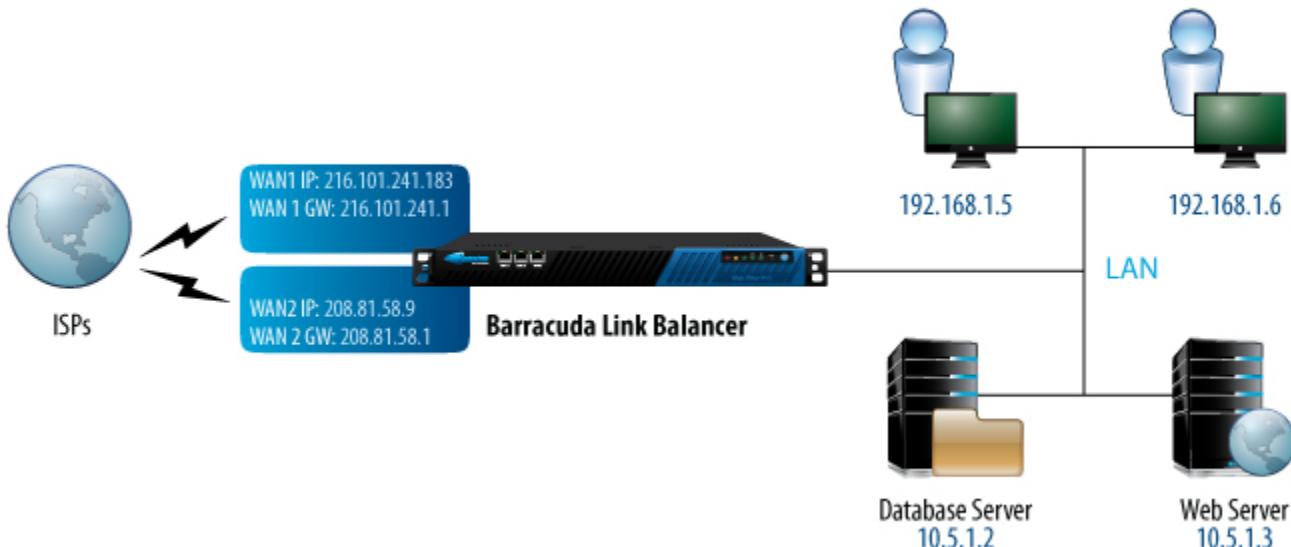
Figure 2.4 shows the example network with a Barracuda Link Balancer installed and acting as a firewall, replacing the customer firewall. A new WAN link has been added.

In this network:

- The Barracuda Link Balancer uses the same IP address for WAN1 that the firewall had used.
- The LAN devices and the LAN interface of the Barracuda Link Balancer must be on a different subnet than all WAN links.
- The Barracuda Link Balancer gateway IP addresses are provided by the ISPs.
- The gateway of the LAN devices is the LAN IP address of the Barracuda Link Balancer.
- Traffic to the servers is passed using port forwarding rules on the Barracuda Link Balancer.

If your servers are externally accessible, reconfigure those servers with private IP addresses. Then create 1:1 NAT rules to map the external IP addresses to the respective private IP addresses of the servers.

Figure 2.4: Deployment Example - Barracuda Link Balancer replacing the firewall



Overview of the Installation Steps

Table 2.2 provides an overview of the steps required to deploy the Barracuda Link Balancer in your network.

Table 2.2: Installation Steps

In Front of Firewall	Replace Your Firewall
Prepare to install, including getting a WAN link with a static IP address.	Prepare to install.
Activate the Barracuda Link Balancer with Temporary Network Settings.	Activate the Barracuda Link Balancer with Temporary Network Settings.
Get Latest Firmware Version.	Get Latest Firmware Version.
Disable the Barracuda Link Balancer Firewall.	Configure Permanent WAN Settings.

Table 2.2: Installation Steps

In Front of Firewall	Replace Your Firewall
Configure WAN and LAN Permanent Settings.	Configure the Barracuda Link Balancer Firewall.
Permanently Install the Barracuda Link Balancer.	Configure Permanent LAN IP Address.
Test Connectivity.	Permanently Install the Barracuda Link Balancer.
	Test Connectivity.

In Front of Your Firewall Installation

These detailed instructions describe how to deploy the Barracuda Link Balancer between the Internet and your firewall. They provide a method to configure the Barracuda Link Balancer completely before connecting it to your production system.



Note: In this mode with the Barracuda Link Balancer's firewall disabled, it is necessary to use an additional static IP address in order to deploy the Barracuda Link Balancer. If you do not have an extra static IP address, you may need to order one from your ISP.

Step 1: Prepare for the Installation

Before installing your Barracuda Link Balancer, complete the following tasks:

1. Verify you have the necessary equipment:

- Barracuda Link Balancer, AC power cord (included)
- Ethernet cables
- a PC with a Web browser

Plug in the Barracuda Link Balancer and power it on.

2. If you are enabling inbound access to resources behind the Barracuda Link Balancer, such as a Web server, you must provide at least one WAN link with a static IP address for receiving the incoming traffic.

Step 2: Activate the Barracuda Link Balancer with Temporary Network Settings

Follow these steps to configure the Barracuda Link Balancer with temporary settings and activate it:

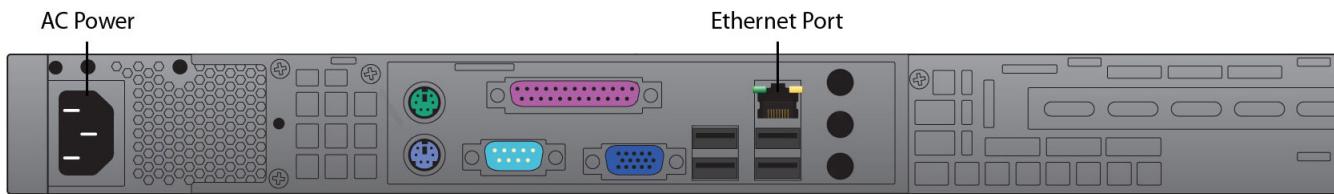
1. Change the network settings of a PC with a Web browser installed to use an IP address of 192.168.200.10, subnet mask of 255.255.255.0 and gateway of 192.168.200.200. Depending on the model, there may be a LAN port on the front of the Barracuda Link Balancer. If there is, connect an Ethernet cable from the PC to that LAN port.

Figure 2.5: Front panel of a Barracuda Link Balancer with LAN port



Otherwise, connect an Ethernet cable from the PC to the LAN port on the back, as shown in [Figure 2.6](#).

Figure 2.6: Back panel of the Barracuda Link Balancer



2. Start the Web browser and access the Web user interface by typing `http://192.168.200.200:8000`. The default username is **admin** and the default password is **admin**.
3. Go to the **Basic > Links** page and double click on one of the WAN ports in the graphic. In the **Links Configuration** section set the **Type** of the WAN link to DHCP to acquire an address in the office network. Alternatively, set the link **Type** to Static and enter a specific IP address. Click **Save Changes**.
4. Connect an Ethernet cable from the corresponding WAN port on the front of the Barracuda Link Balancer into your office network. You should now have Internet connectivity from your PC.
5. At the top of every page, you may see the following warning:

Error: Activation has not been completed. Please activate your Barracuda Link Balancer to enable functionality. ([Click here to activate](#))

Click on the link in the warning message or use the link on the **Basic > Status** page to open up the **Barracuda Networks Product Activation** page in a new browser window. Fill in the required fields and click **Activate**. A confirmation page opens to display the terms of your subscription.

On the **Basic > Status** page, you may need to enter the activation code from the **Barracuda Networks Product Activation** page to activate your Barracuda Link Balancer.



Note: If your subscription status does not change to *Current*, or if you have trouble filling out the **Product Activation** page, call your Barracuda Networks sales representative.

Step 3: Update Firmware

Go to **Advanced > Firmware Update**. If there is a new **Latest General Release** available, perform the following steps to update the system firmware:

1. Read the release notes to learn about the features of this firmware update.
2. Click the **Download Now** button located next to the Latest General Release firmware version. Click **OK** to acknowledge the download duration message. To avoid damaging the Barracuda

Link Balancer, do not power off during an update or download. To view the progress of the download, click **Refresh**. You will be notified when the download is complete.

3. Click **Apply Now** to apply the firmware. Click **OK** to acknowledge the reboot message. Applying the firmware takes a few minutes to complete.
4. After the firmware has been applied, the Barracuda Link Balancer automatically reboots. When the system comes back up, the login page is displayed. Log in again.

Step 4: Disable the Barracuda Link Balancer Firewall

Because you are going to use your existing firewall, you must disable the Barracuda Link Balancer firewall:

1. Go to the **Basic > IP Configuration** page and disable the firewall. Click **OK** to acknowledge the reboot message.
2. The Barracuda Link Balancer will reboot.

Step 5: Configure WAN and LAN Permanent Settings

Configure the permanent settings for the WAN links that will be connected to the Barracuda Link Balancer. Some of the configuration information for the WAN links is provided by your ISP. Be sure to enter these values correctly.

1. Unplug the Ethernet cable connecting the WAN port to your office network.
2. After the Barracuda Link Balancer has rebooted, login to the Web user interface and go to the **Basic > Links** page.
3. For each link that will be connected to this unit:
 - 3a. Click the relevant WAN port in the graphic.
 - 3b. In the Links Configuration section enter the details for the link to be connected to the WAN port.

If the interface uses a static IP address, then you will see the **Additional IP Addresses** list. These are the externally reachable IP addresses that are behind the Barracuda Link Balancer, including the address of your firewall. These need to be identified so that traffic can be accepted and directed to them. The Barracuda Link Balancer is able to locate these addresses automatically, so creating this list is optional but may cause a slight increase in efficiency. To manually create the list, enter the IP addresses or click **Discover** to populate a list of IP addresses of live systems that are on the same Class C network as this WAN interface. The **Discover** button is only visible if there are no entries in the Additional IP Addresses list and if the built-in firewall is disabled.

- 3c. Click **Save Changes**.



Note: If you are enabling inbound access to resources behind the Barracuda Link Balancer, such as a Web server, you must provide at least one WAN link with a static IP address for receiving the incoming traffic.

4. If desired, change the LAN/Management IP address of the Barracuda Link Balancer to its permanent setting. The LAN IP address is only used for management of the Barracuda Link Balancer. (The WAN IP addresses can also be used to access the management interface).

The LAN IP address can be any internal or public address that is reachable through your existing firewall from the LAN. You may allocate an external IP address for it, or choose a non-routable

IP address. If the latter, it should be on a different subnet than the LAN devices already on the network. Remember that if the firewall does not recognize an address as being on the local network it will pass it to the Barracuda Link Balancer.

If the default address of 192.168.200.200 meets this criteria, there is no need to change it.

To change the LAN/Management IP address,

Go to the **Basic > IP Configuration** page and change the Management IP Address and Subnet Mask. Click **Save Changes**. If the address is on a different subnet, your connection will terminate.

5. Power down your Barracuda Link Balancer using the power button on the front of the unit.

Step 6: Install in the Production Network

Now that the Barracuda Link Balancer is configured, install it in its permanent location and connect it to your WAN links:

1. Mount the Barracuda Link Balancer in a 19-inch rack or place it in a stable location. To ensure proper ventilation, do not block the cooling vents on the front and back of the unit.
2. Connect each of the cables from the Internet links into a WAN port on the front of the Barracuda Link Balancer. The ports are labeled WAN1, WAN2, etc. These ports correspond to the WAN ports that you configured in the Web user interface. Be sure to connect them according to your configuration.
3. If there is a LAN port on the front of the Barracuda Link Balancer, connect an Ethernet cable from the outside interface of your existing network firewall to that LAN port. If there is no LAN port on the front, connect the outside interface of your existing network firewall to the LAN Ethernet port on the back panel of the Barracuda Link Balancer. You should see some activity on both the yellow and green lights on the LAN port. If not, you may need to use a crossover cable.

Step 7: Test Connectivity

Now you are ready to test the connectivity to your existing firewall and the systems connected to it. There is no need to change your firewall network configuration - your network firewall should continue to use the ISP provided gateway address.

1. Confirm that you can access the Internet from a client computer on your LAN. If this works, continue.
2. On the test system, log into the Web user interface using the permanent LAN IP address and go to the **Basic > Links** page. The status of each link should appear as **Connected**. You can see the utilization of each link by moving the mouse over the graphic.

On the test system, generate some traffic, by, for example, opening more tabs in the browser of the test system and downloading files from the Internet. You can FTP files from a number of different sites or use torrent to get the traffic to flow on multiple links. Go to the **Basic > Status** page to view graphs that show the incoming and outgoing traffic for each link..



Note: If you have connectivity issues, clear the ARP caches of your existing network components such as the firewall, routers and modems. In some cases, you may need to reboot these devices.

You do not need to update your existing firewall configuration unless you want to make it aware of the new WAN link(s). To do so:

1. Add firewall rules so that traffic from the new links is handled correctly.
2. If you want to be able to manage your existing firewall remotely, add an alias on your firewall for the other links in case the first link is unavailable.

Your Barracuda Link Balancer should be ready for operation. There are a number of other configuration options available. Please refer to the next chapter, *Configuring the Barracuda Link Balancer* on page 19, to review these options.

Replacing Your Firewall Installation

These instructions describe how to deploy the Barracuda Link Balancer as the default gateway for your network, replacing your existing firewall or router.

The steps documented here provide a method to configure the Barracuda Link Balancer completely before connecting it to your production system. This method allows you to copy the firewall configuration of your existing firewall to the Barracuda Link Balancer. A similar process is described in the *Barracuda Link Balancer Quick Start Guide*.

Step 1: Prepare for the Installation

Before installing your Barracuda Link Balancer, verify you have the necessary equipment:

- Barracuda Link Balancer, AC power cord (included)
- Ethernet cables
- a PC with a Web browser

Plug in the Barracuda Link Balancer and power it on.

Step 2: Activate the Barracuda Link Balancer with Temporary Network Settings

Follow these steps to configure the Barracuda Link Balancer with temporary settings and activate it:

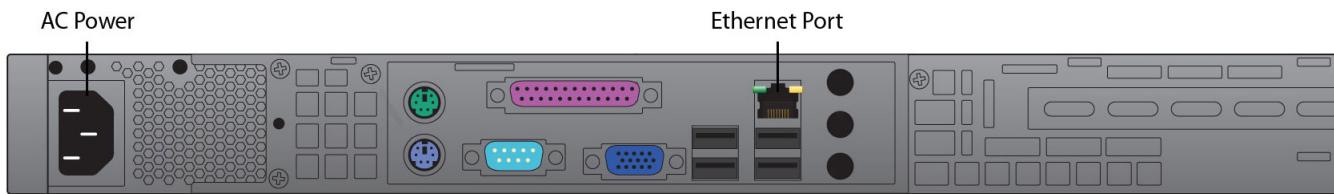
1. Change the network settings of a PC with a Web browser installed to use an IP address of 192.168.200.10, subnet mask of 255.255.255.0 and gateway of 192.168.200.200. Depending on the model, there may be a LAN port on the front of the Barracuda Link Balancer. If there is, connect an Ethernet cable from the PC to that LAN port.

Figure 2.7: Front panel of a Barracuda Link Balancer with LAN port



Otherwise, connect an Ethernet cable from the PC to the LAN port on the back, as shown in [Figure 2.8](#).

Figure 2.8: Back panel of the Barracuda Link Balancer



2. Start the Web browser and access the Web user interface by typing <http://192.168.200.200:8000>. The default username is **admin** and the default password is **admin**.
3. Go to the **Basic > Links** page and double click on one of the WAN ports in the graphic. In the **Links Configuration** section set the **Type** of the WAN link to DHCP to acquire an address in the office network. Alternatively, set the link **Type** to Static and enter a specific IP address. Click **Save Changes**.
4. Connect an Ethernet cable from the corresponding WAN port on the front of the Barracuda Link Balancer into your office network. You should now have Internet connectivity from your PC.
5. At the top of every page, you may see the following warning:

Error: Activation has not been completed. Please activate your Barracuda Link Balancer to enable functionality. ([Click here to activate](#))

Click on the link in the warning message to open up the **Barracuda Networks Product Activation** page in a new browser window. Fill in the required fields and click **Activate**. A confirmation page opens to display the terms of your subscription.

On the **Basic > Status** page, you may need to enter the activation code from the **Barracuda Networks Product Activation** page to activate your Barracuda Link Balancer.



Note: If your subscription status does not change to *Current*, or if you have trouble filling out the **Product Activation** page, call your Barracuda Networks sales representative.

Step 3: Update Firmware

Go to **Advanced > Firmware Update**. If there is a new **Latest General Release** available, perform the following steps to update the system firmware:

1. Read the release notes to learn about the features of this firmware update.
2. Click the **Download Now** button located next to the **Latest General Release** firmware version. Click **OK** to acknowledge the download duration message. To avoid damaging the Barracuda

Link Balancer, do not power off during an update or download. To view the progress of the download, click **Refresh**. You will be notified when the download is complete.

3. Click **Apply Now** to apply the firmware. Click **OK** to acknowledge the reboot message. Applying the firmware takes a few minutes to complete.
4. After the firmware has been applied, the Barracuda Link Balancer automatically reboots. When the system comes back up, the login page is displayed.

Step 4: Configure Permanent WAN Settings

1. After the Barracuda Link Balancer has rebooted, login to the Web user interface and go to the **Basic > Links** page.
2. For each link that will be connected to this unit:
 - 2a. Click the relevant WAN port in the graphic.
 - 2b. In the Links Configuration section enter the details for the link to be connected to the WAN port.
Some of the configuration information is provided by your ISP - be sure to enter it correctly.
 - 2c. Click **Save Changes**.



Note: For each WAN interface that has a static IP address, enter all of your externally accessible servers in the **Additional IP Addresses** field. These need to be identified so that traffic can be accepted for them. You will also need to reconfigure the servers with private IP addresses. In the next step you can create firewall 1:1 NAT rules to direct traffic to those systems.

Step 5: Configure the Barracuda Link Balancer Firewall

Create firewall rules on the Barracuda Link Balancer to match the settings of your current firewall. More information about firewall settings can be found at *Configuring the Firewall on page 21*.

1. Go to the **Firewall > Inbound** page to add inbound rules.
2. Go to the **Firewall > Outbound** page to add outbound rules.
3. Go to the **Firewall > NAT** page to add 1:1 NAT and port forwarding rules.

Step 6: Configure Permanent LAN IP Address

1. Change the LAN IP address of the Barracuda Link Balancer to its permanent setting as the default gateway for your network. To change the LAN IP address, go to the **Basic > IP Configuration** page and change the IP Address and Subnet Mask. Click **Save Changes**.
2. The connection to the PC that you were using will terminate.
3. Unplug the cable connecting your PC to the Barracuda Link Balancer.
4. Power down your Barracuda Link Balancer using the power button on the front of the unit.

Step 7: Install in the Production Network

Now that the Barracuda Link Balancer is configured, install it in its permanent location and make it part of your production network:

1. Mount the Barracuda Link Balancer in a 19-inch rack or place it in a stable location. To ensure proper ventilation, do not block the cooling vents on the front and back of the unit.
2. Connect each of the cables from the Internet links into a WAN port on the front of the Barracuda Link Balancer. The ports are labeled WAN1, WAN2, etc. These ports correspond to the WAN ports that you configured in the Web user interface. Be sure to connect them according to your configuration.
3. Unplug your firewall from the network and plug its LAN connection into the LAN port on the front of the Barracuda Link Balancer, if there is one. If there is no LAN port on the front, plug that connection into the LAN Ethernet port on the back panel of the Barracuda Link Balancer. You should see some activity on both the yellow and green lights on the LAN port. If not, you may need to use a crossover cable.
4. If the IP address of the Barracuda Link Balancer is the same as the IP address of the firewall that you removed, you can ignore this step. Otherwise, make the Barracuda Link Balancer the default gateway for the clients by performing these steps:
 - 4a. Update the configuration of the DHCP server for the clients to give out the LAN IP address of the Barracuda Link Balancer as the default gateway. As the leases are renewed, each client will gain access to the new Internet links.
 - 4b. Change the default gateway of any clients with static IP addresses to the LAN IP address of the Barracuda Link Balancer.

Step 8: Test Connectivity

Test the connectivity to a client system:

1. If needed, change the gateway IP address of a test system to the LAN IP address of the Barracuda Link Balancer.
2. Confirm that you can access the Internet from the test system. If this works, continue.
3. On the test system, log into the Web user interface using the permanent LAN IP address and go to the **Basic > Links** page. The status of each link should appear as **Connected**. You can see the utilization of each link by moving the mouse over the graphic.
4. On the test system, generate some traffic, by, for example, opening more tabs in the browser of the test system and downloading files from the Internet. FTP files from a number of different sites or use BitTorrent to get the traffic to flow on multiple links. Go to the **Basic > Status** page to view graphs that show the incoming and outgoing traffic for each link..



Note: If you have connectivity issues, clear the ARP caches of your existing network components such as routers and modems. In some cases, you may need to reboot these devices.

There are a number of other configuration options available. Please refer to the next chapter, *Configuring the Barracuda Link Balancer* on page 19, to review these options.

Chapter 3

Configuring the Barracuda Link Balancer

This chapter describes the configuration tasks you can perform from the Web user interface after you have completed the installation. The following topics are covered:

<i>Configuring Network Settings</i>	19
<i>Configuring the Firewall</i>	21
<i>Creating Custom Applications</i>	24
<i>Managing Bandwidth</i>	24
<i>Outbound Traffic Routing</i>	25
<i>Configuring Virtual Private Networks</i>	27
<i>Configuring the DNS Server for Inbound Load Balancing</i>	30
<i>Configuring Administrative Settings</i>	35

Configuring Network Settings

This section describes network settings that you can configure using the Web user interface. The following topics are covered:

<i>Adding, Updating or Viewing WAN Link Configuration</i>	19
<i>Adding a New WAN Link</i>	19
<i>Adding Static Routes</i>	20
<i>Configuring VLANs</i>	20
<i>Creating IP Aliases</i>	21
<i>Configuring DNS Servers</i>	21
<i>Configuring Per Interface Health Checks</i>	21
<i>Configuring the DHCP Server</i>	21

Adding, Updating or Viewing WAN Link Configuration

Every WAN link that enters the Barracuda Link Balancer needs to be manually identified. This information is provided by your Internet Service Provider. Use the **Basic > Links** page to identify the links. You can view the configuration information by moving your mouse over the port. To change or add configuration information, click on the corresponding WAN port on the graphic or the expand button in the Configure column.

Adding a New WAN Link

There are a few things to remember if you are adding one or more new WAN links to an already-configured Barracuda Link Balancer. As already mentioned, all links must be configured using the **Basic > Links** page.

New WAN links that are configured correctly are automatically used for outbound link balancing. For inbound traffic, if the Barracuda Link Balancer firewall is enabled, you can add a NAT rule on the **Firewall > NAT** page to map the destination IP address of the traffic on the new link to an internal service.

WAN IP Impersonation

If the Barracuda Link Balancer firewall is disabled, you can avoid having to update rules on your network firewall to include the new WAN link by choosing to map the destination IP address of traffic on the new link to an existing WAN IP address (usually, an address on WAN1). To do this, select the **NAT/Port Forwarding** option on the **Basic > Links** page. Then create a NAT rule on the **Firewall > NAT** page to map the destination IP address of the traffic on the new link to an external IP address on an existing link.

Authoritative DNS

If you are using authoritative DNS to achieve inbound link load balancing, remember to add any new links with static IP addresses to the list of DNS name servers on the **Services > Authoritative DNS** page. See *If You Add a WAN Link After the Domains are Created* on page 34 for more information.

Adding Static Routes

If you have a separate subnet that needs to be able to use the Internet links that are accessible only through the Barracuda Link Balancer, add a static route to specify a gateway for the subnet so that the return traffic can take the correct path. If you have disabled the Barracuda Link Balancer firewall, then static routes can be added to your network firewall. Otherwise, follow these instructions to add static routes to the Barracuda Link Balancer:

1. On the Web user interface, go to **Advanced > Advanced IP Config**. Add the static routes.
2. Test connectivity from each internal network by changing the gateway IP address of a computer on each subnet to the LAN IP address of the Barracuda Link Balancer. Check that you can access the Internet from each subnet.
3. When testing is complete, update the configuration of the DHCP server for the clients to give out the LAN IP address of the Barracuda Link Balancer as the default gateway. As the leases are renewed, each client will have access to all of the new Internet links.
4. Change the default gateway of any clients with static IP addresses to the LAN IP address of the Barracuda Link Balancer.

Configuring VLANs

The Barracuda Link Balancer supports the IEEE 802.1Q standard for explicitly tagging Ethernet frames with VLAN information. Use the **Advanced > Advanced IP Config** page to identify VLANs. Then create a virtual interface that associates an IP address and netmask with a VLAN. Traffic sent to a virtual interface associated with a VLAN will be tagged with the VLAN ID and delivered correctly.

VLANs may not be on the same subnet.

Creating IP Aliases

You can create virtual interfaces or IP aliases by associating an IP address or subnet with a WAN, LAN or VLAN. Each IP address and netmask can only be associated with one WAN, LAN or VLAN.

Virtual interfaces are used:

- to associate an IP address range with a VLAN
- to associate an externally accessible IP address that is on a different subnet than any WAN link with a WAN link.

Create IP aliases using the **Advanced > Advanced IP Config** page.

Configuring DNS Servers

Use the **Basic > Links** page to set the primary and secondary DNS servers for each WAN link. Your ISP should provide you with these settings.

Configuring Per Interface Health Checks

Use the **Basic > Links** page to configure health checks for each link. Multiple methods are supported. You can enter more than one test target (e.g. resolve the DNS domain names of multiple Web sites) to be sure that the link is actually down. Link failure is shown on this page and on the **Basic > Status** page. Also, if a link fails an SNMP trap will be generated, an email will be sent, and an event will be logged.

Configuring the DHCP Server

The Barracuda Link Balancer can act as a DHCP server. Use the **Services > DHCP Server** page to enable the DHCP server and to configure it.

Configuring the Firewall

The Barracuda Link Balancer can act as a firewall, inspecting network traffic as it arrives and allowing or denying passage based on a set of rules. These rules include inbound, outbound, 1:1 NAT and port forwarding rules. Some of this functionality is available even if the firewall is disabled.

This section covers the following topics:

<i>Firewall Functionality</i>	22
<i>Order of Execution of Firewall Rules</i>	22
<i>Inbound Firewall Rules</i>	22
<i>Inbound 1:1 NAT Rules</i>	22
<i>Port Forwarding Rules</i>	23
<i>Outbound Firewall Rules</i>	23
<i>Firewall Logging</i>	23

Firewall Functionality

Using 1:1 NAT and port forwarding rules, the Barracuda Link Balancer can perform:

- 1:1 NAT - Assign external addresses to internal clients.
- Port forwarding (or Port Address Translation) - The traffic to a port across one or multiple links is directed to an internal client.
- Many to 1 NAT - One internal server may receive traffic from more than one WAN link. You can achieve this by creating 1:1 NAT rules or port forwarding rules.
- Port blocking and unblocking.

1:1 NAT and port forwarding rules are executed only if the Barracuda Link Balancer firewall is enabled or, if not, for any WAN link with the **NAT/Port Forwarding** option enabled. Even if the rules are not able to be executed, you can always create rules and save them. This may assist you in configuring the built-in firewall with minimal disruption to your network.

Inbound and outbound firewall rules allow or deny access to remote networks, clients, services and ports. Inbound and outbound firewall rules are executed regardless of firewall status.

The Barracuda Link Balancer firewall also assists in preventing and mitigating distributed denial of service attacks by rate limiting the number of requests that come in to your network.

Order of Execution of Firewall Rules

Firewall rules are arranged in tables from top to bottom in order of precedence. Only the first rule that matches the profile of the traffic is executed.

Inbound Firewall Rules

By default, all connections that are initiated from outside are denied. Add inbound firewall rules to allow exceptions for specific IP addresses, ports and applications. Applications let you define rules that apply to more than one port.

Use the **Firewall > Inbound** page to create firewall rules for incoming packets. If you want to create an inbound rule for an application that is not in the list presented when you add the rule, first go to the **Policy > Applications** page and define a new application.

Inbound 1:1 NAT Rules

When the Barracuda Link Balancer firewall is enabled, externally reachable servers cannot have public IP addresses. You will need to reconfigure these servers with private IP addresses. Identify the public IP addresses as the Additional IP Addresses for a WAN interface that has a static IP address. Then you can create 1:1 NAT rules to direct traffic to your servers.

You can add the public IP addresses as Additional IP Addresses to more than one WAN interface that has a static IP address. All incoming traffic will be forwarded according to the rules you create. This allows traffic to be received by the same internal server from more than one WAN link.

1:1 NAT applies to the IP address only, leaving ports the same on both IP addresses. 1:1 NAT is bi-directional – outbound traffic will include the servers' public IP addresses.

If the Barracuda Link Balancer firewall is disabled, you can create a NAT rule to map the destination IP address of the inbound traffic on one WAN link to another WAN link's IP address. This allows you to add a new WAN link without having to update rules on your network firewall. See *Adding, Updating or Viewing WAN Link Configuration* on page 19 for more details.

When a 1:1 NAT rule is created, an inbound firewall rule to accept traffic for the external IP address is automatically generated. Without this rule, all connections that are initiated from outside are denied. You can view and change this rule – it has a similar Rule Name – using the **Firewall > Inbound** page. You may want to modify that rule to limit access to only those ports or applications that you want to be publicly accessible.

Use the **Firewall > NAT** page to create 1:1 NAT rules and port forwarding rules. If you create a 1:1 NAT rule for an address, there is no need to also create a port forwarding rule.

Port Forwarding Rules

Create port forwarding rules to direct traffic on an external port to a port on an internal IP address. You must specify which WAN link to be used to listen for incoming packets on the port. The return path is handled automatically.

The listen IP address on a specific WAN interface could either be the WAN IP address or any Additional IP address on the same WAN interface. A WAN IP address that is used in any port forwarding rule can not also be used in a 1:1 NAT rule.

You can forward the traffic from a port on multiple WAN links to a port on one internal IP address by creating a rule for each WAN link.

When you add a port forwarding rule, an inbound firewall rule is created automatically to accept traffic on the listen link and port for the private IP address of the server. Without this rule, all connections that are initiated from outside are denied. You can view and change this rule – it has a similar Rule Name – using the **Firewall > Inbound** page.

To add a new port forwarding rule, go to the **Firewall > NAT** page.

Outbound Firewall Rules

By default, all outbound connections are allowed. You can create outbound firewall rules to deny outbound connectivity. For example, you may want to block access to certain online gaming sites that use specific ports.

Use the **Firewall > Outbound** page to create, modify or delete outbound firewall rules. The rules are arranged in the table from top to bottom in order of precedence. Only the first rule that matches the profile of the traffic is executed.

If you want to create an outbound rule for an application that is not in the list presented when you add the rule, go to the **Policy > Applications** page and define a Custom Application.

Firewall Logging

You can view the firewall log displayed on the **Logs > Firewall Log** page to see rules that have been executed and whether the traffic was dropped or allowed. Only rules that have the **Log** check box selected in their rule entry (under the **Firewall** tab) are logged in this way.

Creating Custom Applications

Use the **Policy > Applications** page to view and define applications that can be used in firewall and Quality of Service rules. An application is a combination of a protocol and one or more ports. You can create new applications or use the predefined ones, such as DNS, email, and HTTP.

Managing Bandwidth

The Barracuda Link Balancer allows you to prioritize and control incoming and outgoing traffic and link usage in a variety of ways.

This section covers the following topics:

<i>Link Usage for Inbound and Outbound Traffic</i>	24
<i>Creating Bandwidth or Quality of Service (QoS) Rules</i>	25

Link Usage for Inbound and Outbound Traffic

Outbound Link Balancing

By default, when outbound traffic from a client IP address going to a new destination IP address is detected, the weights of the links are compared and the link with the highest weight is used. The weight of a link is the available capacity based on configured link speed and current usage. The weight for each primary and backup WAN link is calculated on an ongoing and frequent basis.

Inbound Link Balancing

Inbound link balancing and failover are available only if the Barracuda Link Balancer acts as an authoritative DNS server for domains behind it.

When the Barracuda Link Balancer receives a DNS query for a hosted domain, it returns the IP address of a WAN link which the client then uses to reach the hosted domain. The algorithm used to select the returned WAN link is the same as the algorithm used for outbound link balancing.

Grouping WAN Links for Inbound and Outbound Traffic

You can change how frequently WAN links are used by assigning each link to one of three groups for both inbound and outbound traffic. This feature allows you to reserve links for certain types of traffic or to use higher cost links only if the lower cost links fail or become saturated.

The supported usage groups are:

- Primary links — used first for link balancing.
- Backup links — used only if the primary links are down or if they become saturated. If all primary links are down or saturated, then traffic is distributed across all available backup links until the primary links become available.

Private links — used only for traffic that matches IP/Application routing rules or if specified explicitly in the configuration of a VPN. Private links are not used at all for default outbound or inbound link balancing. They are used only if explicitly referenced.

If you want to employ default link balancing policy, where the link with the greatest available capacity is used, set the usage group for each link to **Primary**.

To assign each link a usage group, edit the link on the **Basic > Links** page.

Specifying WAN Link for Outbound Traffic

You can override the link balancing algorithm by creating rules that determine which WAN link certain kinds of outbound traffic will use. See *Specifying the Link Used by Outgoing Traffic* on page 26.

Creating Bandwidth or Quality of Service (QoS) Rules

You can create bandwidth rules that specify the priority that the Barracuda Link Balancer gives to outbound traffic. Bandwidth rules are executed after it has been determined which WAN link the traffic will use. These rules apply to all traffic, including VPN traffic, regardless of whether the Barracuda Link Balancer firewall is enabled or not.

Each rule describes a set of traffic based on one or more parameters: source IP address or range, destination IP address or range, application or applications, time, day of the week, and WAN link. If the conditions are met, the rule assigns a bandwidth priority and a contention priority.

The bandwidth and contention priority classes are compliant with the DiffServ specification. DiffServ bits are inserted if the traffic matches the Quality of Service rules on this page. Then, bandwidth and contention priorities are applied based on those bits, even if they were set by another device. Traffic that does not have any DiffServ bits set is assigned the Default class.

If the amount of traffic on the network is more than what can be sent, traffic from high priority applications is allocated a greater share of the bandwidth. The contention priority subdivides traffic that has the same bandwidth priority level.

Some examples of QoS rules:

- Give higher priority to traffic originating from a set of IP addresses.
- Assign lower priority to FTP traffic so that uploading and downloading of files does not impact other applications.
- Increase the priority of SIP traffic so that calls are not dropped.
- Give VPN traffic a high priority. You can do this by creating a rule where the source IP address is the local VPN endpoint and the destination IP address is the remote VPN endpoint.

Configure QoS rules using the **Policy > Bandwidth Mgmt** page.

Configure the priority classes for the QoS rules on the **Policy > Configuration** page. There are 8 classes per WAN link, each consisting of a minimum and maximum value. Traffic with a given class is guaranteed to get at least the minimum bandwidth and will get up to the maximum if it is available.

Outbound Traffic Routing

By default, all outgoing traffic is link balanced and NAT'd. Also, the source IP address of outgoing traffic is the WAN link that is used by the traffic. You can create outbound routing rules to modify these defaults.

This section covers the following topics:

<i>Specifying the Link Used by Outgoing Traffic</i>	26
<i>Changing the Source IP Address of Outgoing Traffic</i>	27

Specifying the Link Used by Outgoing Traffic

To exempt outgoing traffic from link balancing and/or NAT'ing, create IP/application rules using the **Policy > Outbound Routing** page. IP/application routing rules are based on source IP address, application, and/or destination IP address.

The IP/application routing rules are executed before the link load balancing algorithm. Traffic that matches no rule is both link balanced and NAT'd. These rules are executed regardless of the firewall operating mode.

Examples where IP/application routing rules may be useful include:

- If you are an ISP with externally accessible IP addresses (ARIN networks) behind the Barracuda Link Balancer that are not on the same subnet as your WAN interfaces.
- If you have subnets that you want to exempt from link balancing.
- If you have systems such as mail servers or VPN endpoints that send traffic that must maintain the original source IP address.
- If you have applications that you want to exclude from outgoing link balancing and NAT'ing.

Ping Traffic

To direct ping (ICMP) traffic that originates from behind the Barracuda Link Balancer to use a specific WAN link:

- Create a ping application using the **Policy > Applications** page (select ICMP as the protocol, no port range).
- Create one or more IP/application routing rules for the ping application.

As an example, if WAN1 is a private link to an office and WAN2 is a primary link used for other Internet traffic, make two rules: one that directs ping traffic to the office to use WAN1 and one that allows all other ping traffic to use WAN2. (Remember that private links are only used if the link is explicitly referenced).

VPN and Email Rules

During installation, sample disabled IP/application routing rules are automatically created for outgoing VPN and email traffic to prevent it from being link balanced or NAT'd. To enable those rules, select the WAN link to be used for the traffic.

If you would like to link balance outgoing email or VPN traffic because you have created a way to make that acceptable to the receiver, you can leave the rules in their disabled state or delete them. (For example, you may have created multiple SPF or DNS records for the WAN IP addresses).

Externally Accessible IP Addresses

If you would like to direct traffic from externally accessible IP addresses behind the Barracuda Link Balancer to the WAN link that is on the same subnet, create one or more rules where those addresses are the source IP addresses, link balancing and NAT are turned off, and **Primary Link** is set to **Auto**.

If you have a network where the externally accessible IP addresses (ARIN networks that are not in any WAN subnets) can send their traffic on any WAN link, you can create rules so that traffic

originating from those addresses can go out without being NAT'ed. Depending on how the ISP's routers are set up, traffic from these networks can either be link balanced or be bound to one WAN link. For the latter case, select specific primary and backup links.

Changing the Source IP Address of Outgoing Traffic

To set the source IP address of outgoing traffic to a masquerade IP address, rather than the IP address of the WAN link, create outbound source NAT rules using the **Policy > Outbound Routing** page. Outbound source NAT rules consider source IP address (or range) and, optionally, application and WAN link. If a rule match occurs, the specified external IP address is used as the source IP address of the traffic.

The outbound source NAT rules are executed after the WAN link has been determined by the link load balancing algorithm. They are executed regardless of the firewall operating mode.

The rules are arranged in a table on the **Policy > Outbound Routing** page in order of precedence from top to bottom. Only the first rule that matches the profile of the traffic is executed. If the traffic matches a 1:1 NAT Rule the outbound source NAT rules are ignored.

Configuring Virtual Private Networks

The Barracuda Link Balancer can act as an endpoint in a site-to-site VPN tunnel.

This section covers the following topics:

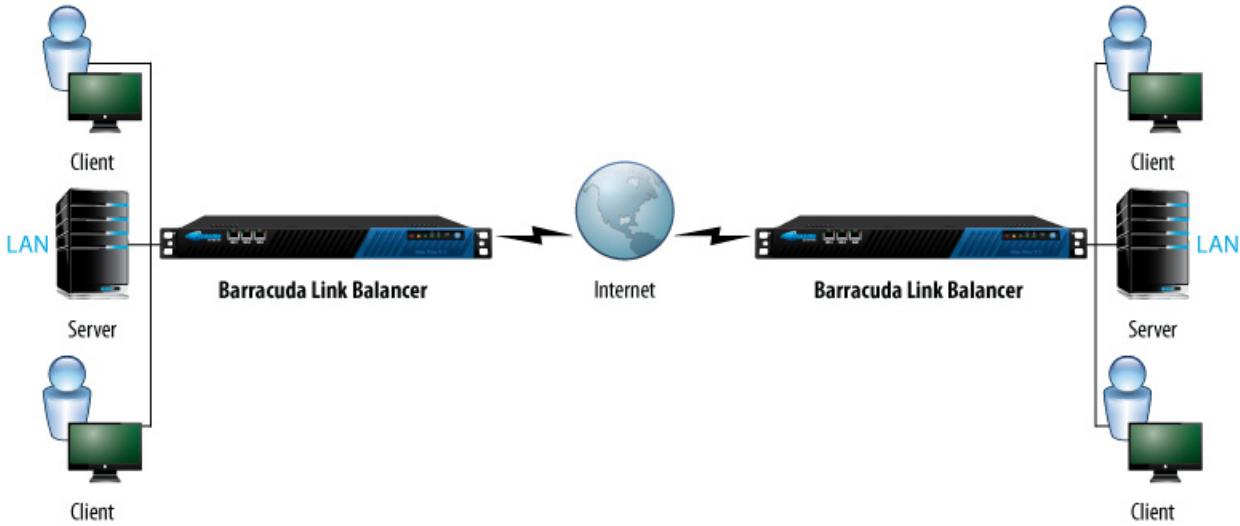
<i>Site-to-Site VPN Tunnels</i>	27
<i>Creating VPN Tunnels</i>	28
<i>Creating a VPN in a NAT'd Environment</i>	28
<i>Failover and Failback</i>	29
<i>VPN Tunnel as Failover Link for a Broken Site-to-Site WAN Link</i> ... 29	
<i>Troubleshooting a VPN Tunnel</i>	29

Site-to-Site VPN Tunnels

You can create a site-to-site VPN tunnel between two Barracuda Link Balancers or between a Barracuda Link Balancer and another device that supports IPsec.

Networks connected via a tunnel will communicate as if they are on the same network, even though they are separated by the Internet.

Figure 3.1: Site-to-Site VPN



The **Services > VPN** page displays all tunnels and their status. You can add, disable, edit or delete a tunnel from this page.

Creating VPN Tunnels

When creating a tunnel, make sure that the relevant tunnel parameters on both ends are in sync. If needed, record the settings on the other endpoint and compare them to the local endpoint. Not matching the settings between the tunnel endpoints is a common cause of failing to establish a tunnel successfully.

Many of the tunnel security parameters are advanced settings and have been given reasonable defaults. If both endpoints are Barracuda Link Balancers use the defaults provided unless you have a specific reason for changing these settings.

For testing purposes, you may choose to start with a shared secret on both endpoints, but using SSL certificates is recommended in a production environment. Upload the local and remote certificates using the **Advanced > Certificates** page.

Creating a VPN in a NAT'd Environment

If either the Barracuda Link Balancer or the remote endpoint is behind a device such as a firewall which is NAT'ing traffic, you must enable the NAT-Traversal (NAT-T) option when creating the VPN tunnel. NAT-T is required to make IPsec and NAT work together. If the option is not enabled, packets will be dropped by the receiving end.

If the remote endpoint for the VPN is behind a NAT'ing device, enter the IP address for the remote endpoint in the **Remote NAT-T IP** field. In this case, the **Primary Remote Gateway** IP address is the NAT'ing device.

If only the local Barracuda Link Balancer is behind a NAT'ing device, the **Primary Remote Gateway** IP address is the remote endpoint and the **Remote NAT-T IP** field should be left blank.

In order for NAT-T to work, open UDP port 4500 on the firewall. The VPN log (on the **Logs > VPN Log** page) will display which VPN endpoint is NAT'd.

Failover and Failback

When configuring a tunnel you can specify a primary and a backup link. If the primary link fails, the tunnel will be reestablished using the backup link. When the primary link is restored, the tunnel will automatically fail back to use the primary link.

VPN Tunnel as Failover Link for a Broken Site-to-Site WAN Link

A VPN tunnel can be configured to act as a failover link replacing a temporarily broken WAN link. To make use of this feature, it is required to have two Barracuda Link Balancers with disabled firewalls in both networks which are to be connected through the failover tunnel. Both Barracuda Link Balancers need to be configured to act as failover WAN endpoints.

To activate the WAN failover, you must select the respective option in the **VPN Status** configuration item of a VPN connection on both Barracuda Link Balancers in order to enable the failover tunnel for WAN1 (or, respectively, one of the other interfaces). If the WAN link fails, the VPN connection will then be activated. When the WAN link is restored, the VPN connection will no longer be used.



Note: External firewalls must be configured properly to allow the VPN failover tunnel.

To make use of this feature, please perform the following configuration tasks:

1. Add an IP/APP rule to send all site-to-site traffic via the WAN link and use the VPN as failover for this traffic.
2. Add an IP/APP rule to send all remaining traffic via any WAN link but not expect this traffic to failover to the VPN.

IP/APP rules should be configured as described below to allow this to happen:

- IP/APP rule #1: Src 192.168.17.0/24, App *, Dst 172.16.1.0/24, LB No, use MPLS, no Backup, no NAT
- IP/APP rule #2: Src 192.168.17.0/24, App Ping, Dst 172.16.1.0/24, LB No, use MPLS, no Backup, no NAT
- IP/APP rule #3: Src 0.0.0.0/0, App *, Dst 0.0.0.0/0, LB No, use DSL, no Backup, NAT yes
- IP/APP rule #4: Src 0.0.0.0/0, App Ping, Dst 0.0.0.0/0, LB No, use DSL, no Backup, NAT yes

Troubleshooting a VPN Tunnel

If the Barracuda Link Balancer is unable to establish a tunnel then you may be able to discover the problem by checking the following:

- On the **Logs > VPN Log** page, check the VPN Log to see if anything has been logged about the cause of the failure.
- On the **Services > VPN** page, click **Edit** next to the tunnel entry to view the tunnel parameters. Check that the security and authentication values match the tunnel parameters of the other end of the tunnel.
- Check the link status using the **Basic > Status** page.

- Use the tools on the **Advanced > Troubleshooting** page, ping the remote gateway and perform other diagnostics on the network connection.

Configuring the DNS Server for Inbound Load Balancing

Configure the Barracuda Link Balancer as an authoritative DNS Server for your domain or domains to achieve inbound link load balancing.

This section covers the following topics:

<i>Introduction</i>	30
<i>DNS Records Time to Live</i>	30
<i>Recommended Deployment</i>	31
<i>Split DNS</i>	31
<i>DNS Zone Transfer Blocking</i>	31
<i>Becoming an Authoritative DNS Host</i>	31
<i>If You Add a WAN Link After the Domains are Created</i>	34
<i>Zones and Domains</i>	34
<i>DNS Records</i>	35

Introduction

The Barracuda Link Balancer can act as an authoritative DNS server, returning definitive answers to DNS queries about domain names installed in its configuration. This allows you to define one or more domains that are accessible via more than one WAN link. When asked to resolve a host, the Barracuda Link Balancer will return one of the IP addresses of the available WAN links. This provides two benefits:

- Failover - If one WAN link goes down, the domain is still available via one of the other WAN links.
- Incoming link balancing - Incoming traffic to the domain will be spread across all links that you configure for that domain.

Only WAN links with static IP addresses can be advertised to respond to DNS queries. However, you can accept traffic on any of your WAN links for a domain configured on the Barracuda Link Balancer.

DNS resource records describe the hosts and name servers and other attributes of the domain. Following the instructions provided here and using the user interface of the Barracuda Link Balancer, you can create the records that describe the domain or domains that are hosted on the LAN side of the Barracuda Link Balancer. The supported DNS resource records are described in the section *DNS Records* on page 35

DNS Records Time to Live

As already mentioned, making the Barracuda Link Balancer the authoritative DNS server for the domains that are behind it increases the availability of your hosted servers. When asked for the IP address of a hostname, the Barracuda Link Balancer returns a DNS A record that contains the IP address of one of your WAN links. Every DNS record has a TTL (Time to Live) value. TTL is the length of time that the DNS record may be cached. For most DNS records, two days is a typical and acceptable value. However, A records should have a very short TTL - we recommend 30 seconds. If a WAN link fails, its address will no longer be returned, so the inbound traffic to this host will not be

disrupted. A short TTL value for this record ensures that the cached address for the failed link times out quickly.

Specifying a short TTL for A records also assists in link balancing. Because the address for a host that is returned varies among the available links, the short TTL guarantees that the link used for incoming traffic directed to that host also varies frequently.

Recommended Deployment

We recommend that you take advantage of this feature if you are hosting services such as Web servers, VPNs and email that are name based. This increases the availability of your services and provides a way to do inbound link balancing.

Split DNS

The Barracuda Link Balancer supports a split DNS infrastructure. If the same hostname is used for a resource that is accessible both internally and externally, internal network clients receive the internal IP address and external clients receive the external IP address when they ask for the address of that hostname. Specifically, the A record for the hostname includes two views, one with the internal IP address and one with the external IP address. In this way, clients only see the address that they should use.

Details of how to make this work can be found in *Step 3 - Set up DNS for Internal Clients* on page 33.

The split DNS infrastructure handles accessing resources using a hostname. What about accessing externally accessible resources using an IP address? If local clients use external IP addresses to access internal servers, the Barracuda Link Balancer translates the address and properly forwards those requests back to internal servers.

DNS Zone Transfer Blocking

The Barracuda Link Balancer can be configured to block zone transfers on some or all of the domains that it hosts. An AXFR/IXFR query that is sent from another DNS server to the Barracuda Link Balancer (to request a copy of the DNS records) is rejected if zone transfers are disabled for that domain. By default, zone transfers are enabled for all domains created.

Becoming an Authoritative DNS Host

Table 3.1 provides an overview of the steps required to make the Barracuda Link Balancer an Authoritative DNS host.

Table 3.1: Configuration Steps

Step	Explanation
Enable authoritative DNS on the Barracuda Link Balancer.	Identify which WAN links are to be used as Name Servers.

Table 3.1: Configuration Steps

Step	Explanation
Create the Domain	Define one or more domains on the Barracuda Link Balancer.
Set up and test DNS for Internal Clients.	Make your internal DNS server forward queries to the Barracuda Link Balancer. Configure split DNS.
Add DNS records.	For Web servers and email servers.
Delegate your domain to the Barracuda Link Balancer from your registration service.	Tell the Internet that your domain exists or has changed.
Test external access.	

Step 1 - Enable Authoritative DNS on the Barracuda Link Balancer

Go to the **Services > Authoritative DNS** page and enable Authoritative DNS and each of the WAN links in the table of DNS Server Listen Links. This table includes all WAN links with static IP addresses (configured on the **Basic > Links** page). You can change the value for the Name Server for each link or keep the default. The Name Server value is used as a label for NS records for all the domains. Enter an unqualified name, for example, ns1

Step 2 - Create one or more Domains

Check that the value for Default Domain specified on the **Basic > IP Configuration** page is accurate.

If the built-in firewall is enabled, and if you have created 1:1 NAT rules and/or port forwarding rules, make sure that they use the correct hostname. You can look at those rules on the **Firewall > Authoritative DNS** page.

On the **Services > Authoritative DNS** page create the domain. When you have done this, you should see that the following records are created:

- Start of Authority (SOA) record
- Name Server (NS) record. One NS record for each name server in the DNS Server Listen Links table is generated.
- Address (A) record - One A record is created for each name server in the DNS Server Listen Links table. An A record is also created for each matching hostname found in 1:1 NAT and Port Forwarding rules, as described in the next section.

If the Barracuda Link Balancer has the firewall enabled:

- When you create a new domain, the Barracuda Link Balancer looks for existing 1:1 NAT and port forwarding rules that include names in the Hostname field that have a domain suffix that is the same as the newly created domain name.
- Or, if you create a domain that is the same as your default domain (as specified on the **Basic > IP Configuration** page), the Barracuda Link Balancer looks for rules that have hostnames that do not appear to be fully qualified domain names.
- In either case, an A record for each matching rule, including both external and internal addresses, will be automatically created for each hostname.

The DNS records are created with typical default values. You can see all of the values for each record and change them by clicking **Edit** next to the record in the DNS Records section.

Step 3 - Set up DNS for Internal Clients

If you have an internal DNS server, configure it to forward queries to the LAN IP address of the Barracuda Link Balancer.

If the built-in firewall of the Barracuda Link Balancer is enabled:

As already described, when you create a new domain, the Barracuda Link Balancer looks for existing 1:1 NAT and port forwarding rules that include names in the Hostname field that have appear to be relevant and creates an A record for each matching rule, including both external and internal addresses.

In some cases, this mapping will not reflect your configuration. Using an internal network client, try to access a hostname for a resource that is available both internally and externally. If the test fails, edit the A record for the unresolved hostname. The **DNS Record** page will appear. In the IP Addresses table, add addresses to the Local Network column to be used in response to internal DNS queries.

If the built-in firewall of the Barracuda Link Balancer is disabled:

The Barracuda is not able to derive internal a mapping between external and internal addresses if the firewall is disabled. If you want internal addresses to be served, edit the A record for the hostname of each resource that is available both internally and externally. The **DNS Record** page will appear. In the IP Addresses table, add addresses to the Local Network column to be used in response to internal DNS queries.

Using an internal network client, test your changes by trying to access the resource using its hostname.

Step 4 - Add More DNS Records

Add more DNS records to your domain(s) to match your configuration. For example, each email server needs an MX record and a corresponding A record. Each Web server needs an A record.

If you have externally reachable IP addresses that are not tied to any interface, such as ARIN networks, create an A record for each one:

- If the address is not routed through the Barracuda Link Balancer, select CUSTOM in the **Links** list.
- If the address is routed through the Barracuda Link Balancer, select ANY in the **Links** list.

Step 5 - Update Your Domain Registrar

If you haven't already registered your domain name, register it with a domain name registrar like GoDaddy.com or register.com. Make the NS records of the domain point to your static WAN IP addresses.

If your domain name is already registered, contact your registrar to update the NS records of the domain to point to your static WAN IP addresses. Remove records that reference the domain or domains that are now delegated to the Barracuda Link Balancer.

Hosting a sub-domain

If your domain is hosted at your ISP or elsewhere and you want to delegate a sub-domain to be resolved by the Barracuda Link Balancer, you will have to add some records to the zone file of the domain where it is stored at the registrar. If the domain is example.com, and you want to host my.example.com and you have two name servers ns1 and ns2, add these lines, using the actual IP addresses of your name servers:

```
my    IN  NS   ns1
my    IN  NS   ns2
ns1   IN  A   216.101.241.181
ns2   IN  A   192.0.2.2
```

Then you can create the my.example.com. domain on the Barracuda Link Balancer.

Step 6 - Test

From a host on the Internet, run nslookup on your domain name(s). The returned IP addresses should be the IP addresses of your WAN listen links.

Depending on the change, it may take some time for your changes to be noted throughout the Internet, depending on how long the various resolvers cache DNS responses. For example, it may take a day before a new domain name is accessible via the Internet. If a domain name was previously registered and the DNS record is modified, any server on the Internet that has the previous information will not get the update until the TTL of the original record has passed.

If You Add a WAN Link After the Domains are Created

If, after creating your domains, you add a new WAN link, complete these steps to use the new link for DNS queries (static links only) and inbound link balancing:

1. Go to the **Services > Authoritative DNS** page.
2. If this is a static link and you want it to be used to respond to DNS queries:
 - Identify the new link as a DNS Server Listen Link and assign it a Name Server label.
 - For each domain that is already defined, add a new NS record and a new A record to each domain for the new link.
3. Edit the A records for your servers to enable inbound traffic to be received on the new link for the corresponding internal servers. Specifically, when you edit the A record, on the **DNS Record** page you can select the new WAN link from the **Links** list and add it to the A record.

Zones and Domains

A domain name server stores information about part of the domain name space called a zone. All names in a given zone share the same domain suffix. For example, if barracuda.com is the domain suffix, mail.barracuda.com and eng.barracuda.com are possible sub-domains. These may be all served by one domain name server or some of the sub-domains may be delegated to other domain name servers. Every domain or sub-domain is in exactly one zone.

Rather than make a distinction between a zone and a domain, the user interface of the Barracuda Link Balancer simply asks you to create a domain.

DNS Records

DNS Records Generated when Creating a Domain

When you create a domain on the Barracuda Link Balancer the following records are automatically generated:

- Start of Authority (SOA) record - The SOA record defines the global parameters for the hosted domain or zone. Only one SOA record is allowed per hosted domain or zone.
- Name Server (NS) record - NS records specify the authoritative name servers for this domain. One NS record for each name server in the DNS Server Listen Links table is generated.
- Address (A) record - A records map a hostname to an IP address. Each host inside the domain should be represented by an A record. One A record is created for each name server in the DNS Server Listen Links table. An A record is also created for each matching domain name found in 1:1 NAT and Port Forwarding rules.

Additional DNS Records

Once a zone has been created, you can edit the above records or add NS, A and any of the following records to a zone:

- Mail Exchanger (MX) record - MX records point to the email servers that are responsible for handling email for a given domain. There should be an MX record for each email server, including backup email servers if they exist. If an email server lies within the domain it requires an A record for each name server. If the email server is outside the domain, specify the FQDN of the server, ending with a dot. Example: mail.my-isp.net.
- Text (TXT) record - Text records allow text to be associated with a name. This can be used to specify Sender Policy Framework (SPF) or DomainKeys records for the domain.
- Canonical Name (CNAME) record - A CNAME record provides a mapping between this alias and the true, or canonical, hostname of the computer. It is commonly used to hide changes to the internal DNS structure. External users can use an unchanging alias while the internal names are updated. If the real server is outside the domain, specify the FQDN of the server, ending with a dot. Example: server1.my-isp.net. If a domain name has a CNAME record associated with it, then it can not have any other record types. Do not use CNAME defined hostnames in MX records.
- Service (SRV) record - Service records are used to store the location of newer protocols, such as SIP, LDAP, IMAP and HTTP.
- Pointer (PTR) record - PTR records point to a canonical name. The most common use is to provide a way to associate a domain name with an IP address.
- Other (OTHER) record - Use an OTHER record to add a type of DNS record that is not supported, such as NAPTR.

More information about these records and their attributes can be found in the online help.

Configuring Administrative Settings

This section describes the configuration tasks you can perform from the Web user interface. The following topics are covered:

<i>Controlling Access to the Web User Interface</i>	36
<i>Changing the Default Password</i>	36

<i>Setting Email Addresses for Alerts</i>	36
<i>Customizing the Appearance of the Web User Interface</i>	36
<i>Setting the Time Zone of the System</i>	36
<i>Enabling SSL for Administration</i>	37

Controlling Access to the Web User Interface

To control access to the Web user interface, navigate to the **Basic > Administration** page. It allows you to perform the following tasks:

- Allow or deny administration access using the WAN interfaces. Denying access from the WAN interfaces is one way to prevent brute force login attacks on your system. You cannot disable administration access via the LAN.
- Specify the IP addresses or subnet masks of the systems that can access the Web user interface. Attempts to log in from other systems will be denied.
- Change the HTTP port used to access the Web user interface (default is port 8000).
- Change the length of time of inactivity allowed until the administrator is logged out of the Web user interface.

Changing the Default Password

To prevent unauthorized use, change the default administrator password for the Web user interface to a more secure password using the **Basic > Administration** page.

Setting Email Addresses for Alerts

Alert emails are generated automatically by the Barracuda Link Balancer to notify you when, for example, a link is down or if your system is low on disk space. Every SNMP trap (except for the WANx saturated trap) generated causes an email to be sent. Specify the email address that is sent alerts from the Barracuda Link Balancer using the **Basic > Administration** page. To enter multiple addresses, separate each address with a comma. Alert emails, if any have been generated, are sent hourly.

On the **Basic > IP Configuration** page, enter the default hostname and default domain name of the Barracuda Link Balancer. The default hostname and the default domain name are displayed in all alert emails sent by the Barracuda Link Balancer.

Customizing the Appearance of the Web User Interface

Use the **Advanced > Appearance** page to customize the default images used on the Web user interface. This tab is only displayed on certain Barracuda Link Balancer models.

Setting the Time Zone of the System

Use the **Basic > Administration** page to set the time zone of your Barracuda Link Balancer. The current time on the system is automatically updated via Network Time Protocol (NTP).

It is important that the time zone is set correctly because this information is used to coordinate traffic distribution and in all logs and reports. If two Barracuda Link Balancers are to be clustered, the time zone must be the same on both before the cluster can be created.



Note: The Barracuda Link Balancer automatically reboots when you change the timezone.

Enabling SSL for Administration

You can choose to require that only secure SSL connections can access the Web user interface. SSL ensures that your passwords and the rest of the data transmitted to and received from the Web user interface are encrypted. The **Advanced > Secure Administration** page allows you to configure SSL.

In order to only allow secured connections when accessing the Web user interface, you need to supply a digital SSL certificate which will be stored on the Barracuda Link Balancer. This certificate is used as part of the connection process between client and server (in this case, a browser and the Web user interface on the Barracuda Link Balancer). The certificate contains the server name, the trusted certificate authority, and the server's public encryption key.

The SSL certificate which you supply may be either private or trusted. A private, or self-signed, certificate provides strong encryption without the cost of purchasing a certificate from a trusted certificate authority (CA). However, the client Web browser will be unable to verify the authenticity of the certificate and a warning will be sent about the unverified certificate. To avoid this warning, download the Private Root Certificate and import it into each browser that accesses the Barracuda Link Balancer Web user interface. You may create your own private certificate using the **Advanced > Secure Administration** page.

Instead of a private certification, you may also use the default pre-loaded Barracuda Networks certificate. The client Web browser will display a warning because the hostname of this certificate is "barracuda.barracudanetworks.com" and it is not a trusted certificate. Because of this, access to the Web user interface using the default certificate may be less secure.

A trusted certificate is a certificate signed by a trusted certificate authority (CA). The benefit of this certificate type is that the signed certificate is recognized by the browser as trusted, thus preventing the need for manual download of the Private Root Certificate. Use the **Advanced > Secure Administration** page to create a Certificate Signing Request which you can submit to a Certificate Authority to purchase a trusted certificate.

Creating a High Availability Environment

This chapter describes how to create a high availability environment by clustering two Barracuda Link Balancers. It includes the following topics:

<i>Overview</i>	39
<i>Planning Your High Availability Deployment</i>	42
<i>Creating a Cluster</i>	44
<i>Removing a System from a Cluster</i>	46
<i>Updating Firmware on Clustered Systems</i>	46

Overview

The High Availability option allows you to link two Barracuda Link Balancers as a clustered active-passive pair. Both systems are connected to the WAN links, but only one is actively processing traffic at any time. The two systems continuously share almost all configuration settings and monitor each other's health.

If clustering two Barracuda Link Balancers is not a viable option, as an alternative, consider configuring Ethernet Passthrough. This feature is only available on certain models.

Ethernet Passthrough

If Ethernet Passthrough is configured and if the Barracuda Link Balancer fails, all traffic from WAN1 will be passed directly to the LAN. Do not enable this feature in these cases:

- If your network is relying on the Barracuda Link Balancer firewall to perform IP or port address translation for internal IP addresses.
- If you have clustered systems, because the passive system will take over if this system fails.

Configure the Ethernet Passthrough option by using the **Advanced > High Availability** page.

Operation of High Availability (HA)

The active system in a clustered pair handles all of the traffic until one of the following components experiences a failure or an outage:

- Its connection to the LAN.
- All of its WAN links (administrator configurable option).
- The Barracuda Link Balancer itself.

When one of these conditions is detected, the passive system becomes active and link balances the traffic from the WAN links.

Clustered Barracuda Link Balancers communicate according to the Virtual Router Redundancy Protocol (VRRP) specification. Both are configured with a single virtual IP address called the VRRP virtual IP address. This address is serviced only by the active system. If the Barracuda Link Balancer firewall is enabled, then the VRRP virtual IP address is the default gateway for devices on the LAN.

In the event of a system failure, the other system in the cluster will assume the VRRP virtual IP address and take on the role of the active system in the cluster. An alert message will be sent to the administrator.

It is recommended that you use the VRRP virtual IP address to manage the Barracuda Link Balancer since that always points to the active system. Changes will automatically be propagated to the passive system.

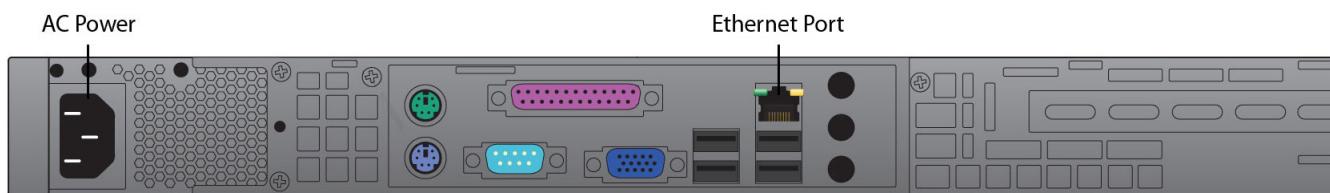
Physical Connectivity of the Clustered Systems

All Barracuda Link Balancer cluster pairs may be linked using the LAN interface. Certain models also support a LAN2 interface: if there is a physical LAN port on the front panel, the Ethernet port on the back is the LAN2 port.

Figure 4.1: Front panel of a Barracuda Link Balancer with LAN port



Figure 4.2: Back panel of the Barracuda Link Balancer



Linking the two systems using the LAN2 port ensures that communication between the two systems will not be delayed or compromised by other traffic on the LAN. It increases the reliability of the connection between the two systems and may reduce the time required for failover to occur.

Use a crossover cable between the LAN2 ports to connect the two systems. The LAN2 IP addresses must be on the same subnet.

Requirements for Clustered Systems

Before joining two systems together, each Barracuda Link Balancer must meet the following requirements:

- Be model 330 or higher.
- Be the same model as the other Barracuda Link Balancer.

- Be activated and on the same version of firmware. The High Availability capability is only available on firmware 2.x and later.
- Be able to reach the other Barracuda Link Balancer on the LAN interface. This last requirement applies only if you do not plan to use the LAN2 port for clustering.

Synchronization of Data Between Clustered Systems

When two Barracuda Link Balancers are initially joined, most configuration data, such as WAN settings, firewall rules, VPN settings and operating mode, is copied from the primary system in the cluster to the backup system (the system that joins the cluster). This configuration data is synchronized between the systems on an ongoing basis.

However, these configuration data are unique and are *not* synchronized between the two systems:

- LAN IP address, LAN2 IP address, DNS servers, default domain and time zone.
- System password, time zone and Web interface HTTP port, as configured on the **Basic > Administration** page.
- All parameters on the **Advanced > Appearance** page.
- The HTTPS port and SSL certificate used to access the Web interface, as configured on the **Advanced > Secure Administration** page.

Failover and Fallback

There is an automatic failback option that can be configured if you want the originally active (primary) system to resume link balancing upon its recovery after a failover. This option can be found on the **Advanced > High Availability** page. Alternatively, you can manually switch to the primary system using the Failback command that is available on the same page.

Planning Your High Availability Deployment

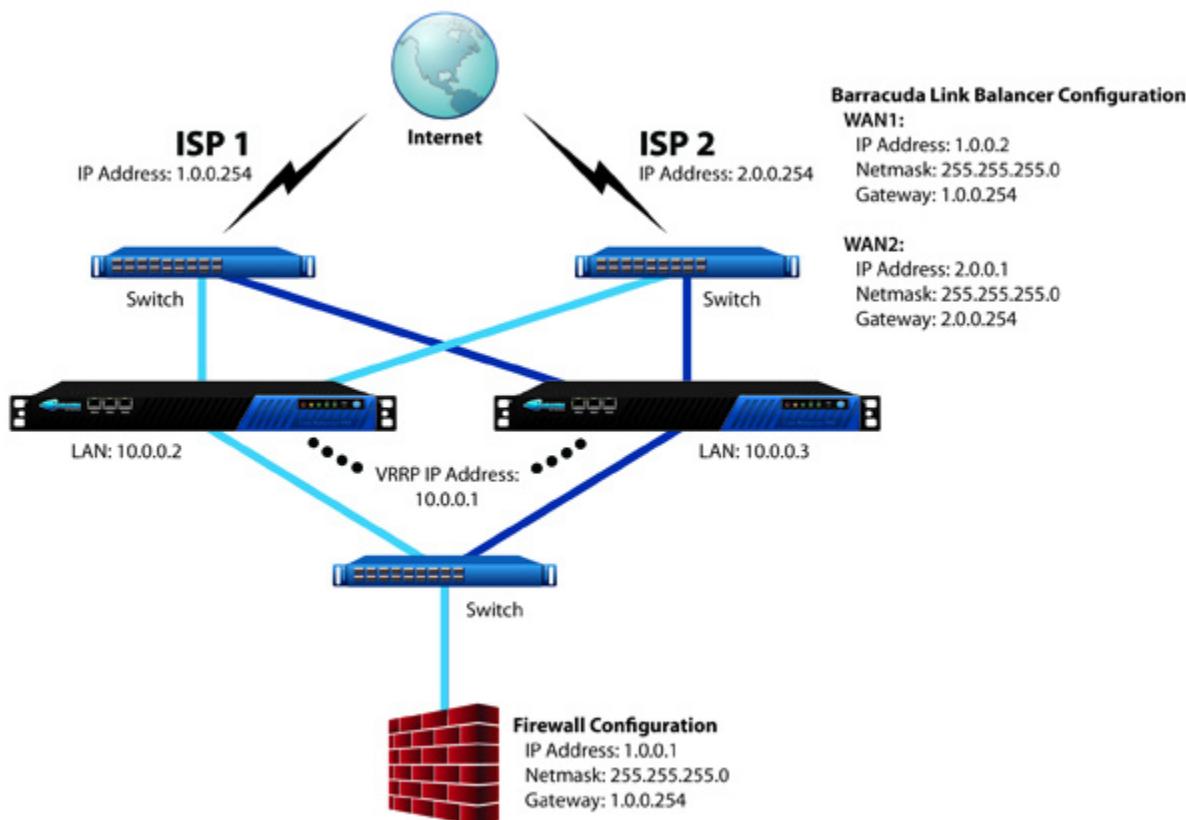
Extra equipment may be needed to support the clustered Barracuda Link Balancers. You may need to add switches so that the WAN links can connect to two systems. If deploying in front of an existing firewall, you will need to add a switch between the Barracuda Link Balancers and the firewall (or two switches for dual firewalls).

The following figures show examples of deployments of a pair of clustered Barracuda Link Balancers with two clustered firewalls, with one firewall and with no external firewall.

In Front of Single Network Firewall

Figure 4.3 shows two Barracuda Link Balancers deployed with one network firewall. The LAN IP addresses of the two Barracuda Link Balancers and the VRRP virtual IP address must all be on the same subnet.

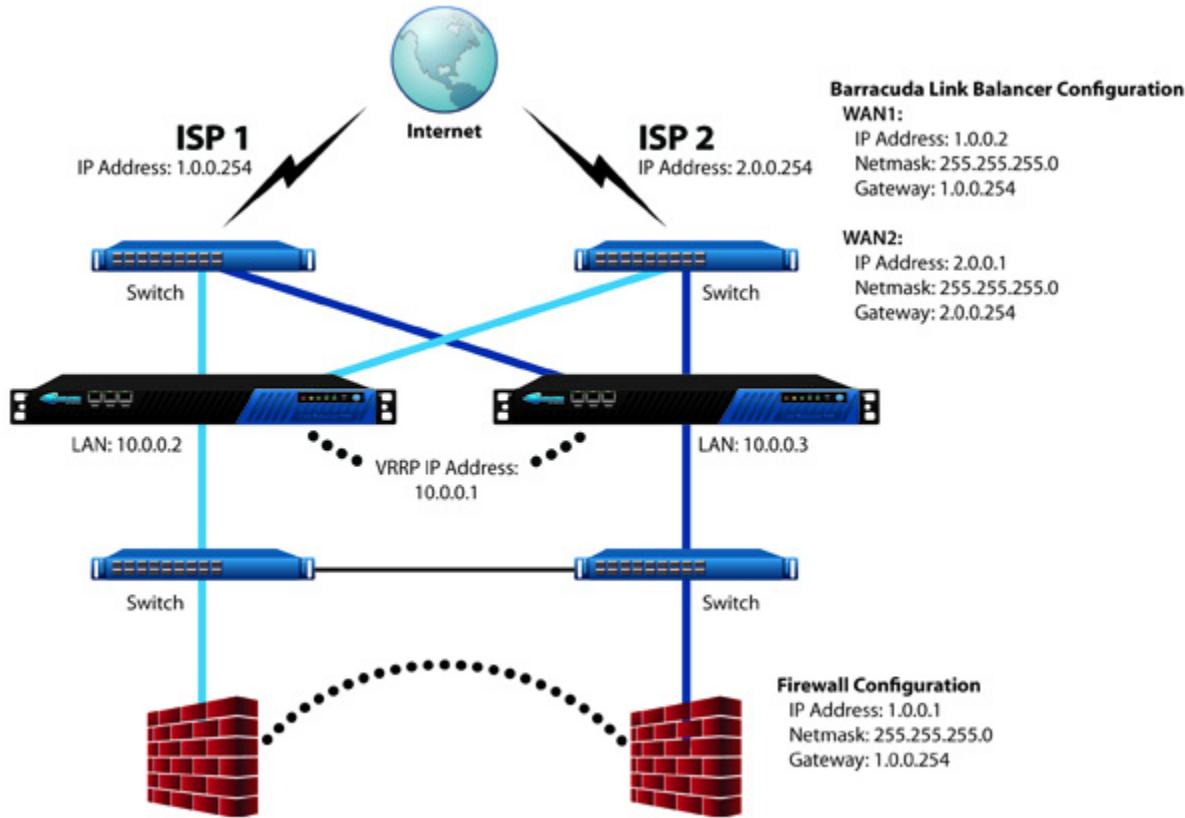
Figure 4.3: Deployment Example - Single Firewall



In Front of Dual Network Firewalls

Figure 4.4 shows two Barracuda Link Balancers and two clustered firewalls. The LAN IP addresses of the two Barracuda Link Balancers and the VRRP virtual IP address must all be on the same subnet.

Figure 4.4: Deployment Example - Dual Firewalls

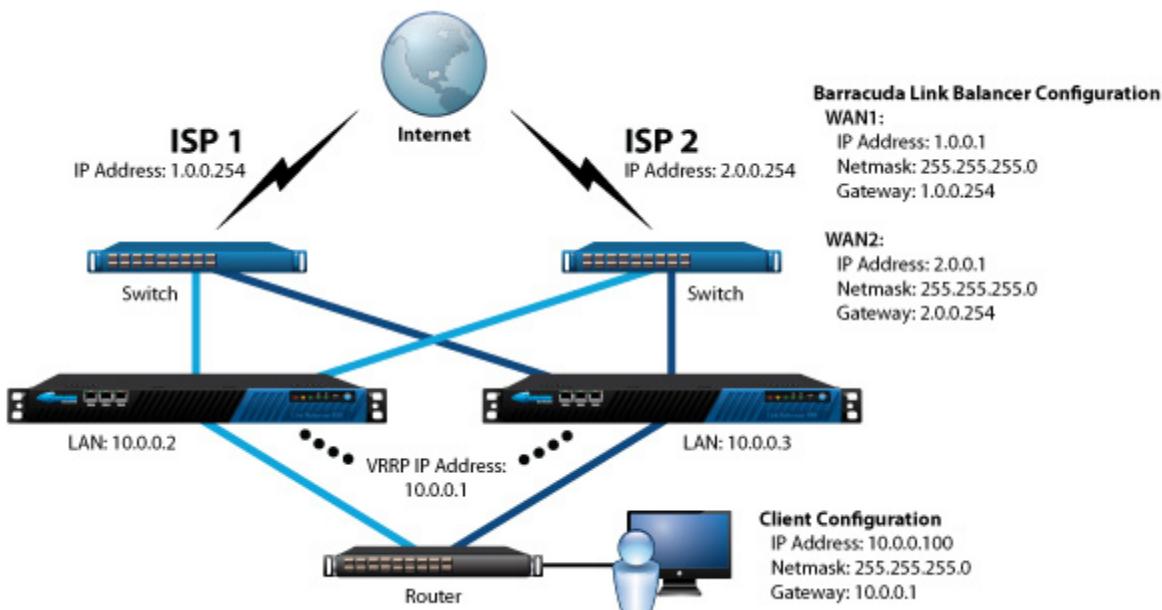


No External Firewalls

Figure 4.5 shows two Barracuda Link Balancers with the firewall enabled. As in the other deployment examples, the LAN IP addresses of the two Barracuda Link Balancers and the VRRP virtual IP address must all be on the same subnet. Note that only in this example the VRRP virtual IP address is the default gateway for devices on the LAN.

If you are adding a second Barracuda Link Balancer to a network where the gateway of the client devices were already configured to use the LAN IP address of the first Barracuda Link Balancer, you can assign a new LAN IP address to that Barracuda Link Balancer, and use its original LAN IP address as the VRRP virtual IP address.

Figure 4.5: Deployment Example - No External Firewall



Creating a Cluster

These instructions describe how to deploy a pair of Barracuda Link Balancers in a cluster. You can also find detailed configuration steps and descriptions of clustering parameters in the online help of the Web interface.

Prior to clustering the Barracuda Link Balancers, you should first place the systems into their production IP address range. This will allow you to avoid having to reconfigure the cluster later because of an IP address change.

Step 1: Complete Installation Process for Both Systems

To prepare the Barracuda Link Balancers for clustering, put both systems in the production location on the network. Complete the following steps:

1. If the primary system is a brand new, unconfigured system, then you will need to completely install, configure and test the primary system as described in *Installing the Barracuda Link Balancer* on page 7.

If the primary system is already configured and operational, update its firmware. If the firewall is enabled, change its LAN IP address to a new value. The original LAN IP address will be used as the VRRP virtual IP address.

2. On the backup system:

- Configure a WAN link and connect to the Internet.
- Activate the system (**Basic > Status**).
- Configure the LAN IP address, LAN2 IP address (optional) and the default domain (**Basic > IP Configuration**).
- Set the time zone to be the same as that of the primary system (**Basic > Administration**).
- Update the firmware (**Advanced > Firmware Update**).
- Connect to the LAN.
- If using the LAN2 ports, connect them using a crossover cable.

Step 2: Create the Cluster

Navigate to the **Advanced > High Availability** page on both systems and perform the following steps:

1. On the primary system, enter values into the fields in the **Cluster Settings** section, and save your changes.
2. On the backup system, enter the same values for the cluster settings and save your changes. In the **Clustered Systems** section, enter the LAN2 IP address of the primary system (if it is to be used) or the LAN IP address of the primary system and click **Join Cluster**. The backup system will reboot.
3. When the backup system comes back up, refresh the **Advanced > High Availability** page on both systems and verify that
 - Each system's LAN or LAN2 IP address appears in the **Clustered Systems** table.
 - The status of each system is green.

The systems are now joined. The shared configuration settings of the primary system (as listed in *Synchronization of Data Between Clustered Systems* on page 41) will be copied to the secondary system, and each system will begin monitoring the health of the other system.

Step 3: Connect WAN Links and Test

Connect all WAN links to the backup Barracuda Link Balancer.

To test the clustering capability, power off the active system. The backup system should take over the link balancing function. Power the system back on to test automatic failback.

If you chose the option to failover if all WAN links are down, then test this by removing the WAN links from the active system.

Step 4: Put in Production

From now on, always use the VRRP virtual IP address to manage the Barracuda Link Balancer so that you can be sure that any changes that you make will occur immediately on the active system.

Removing a System from a Cluster

You can find detailed instructions to separate clustered Barracuda Link Balancers in the online help. The system that is to be removed from the cluster must have its cluster settings erased and its network links disconnected. If another system is not taking its place in the cluster, the remaining system should have its cluster settings removed, and possibly also have its LAN IP address changed to what had been the VRRP virtual IP address.

Use the LAN IP addresses of the Barracuda Link Balancers to access their Web interfaces while separating them. As soon as one system is removed from the cluster the VRRP virtual IP address will not be usable.

Updating Firmware on Clustered Systems

In order to ensure that the transfer of control from one system to the other does not happen while a firmware update is in process, you must disable automatic failback on the **Advanced > High Availability** page before attempting to update the firmware on either system. We recommend that you update the firmware on the passive system first, and then update the firmware on the active system. Enable automatic failback again, if desired, after both systems have been updated and are back online.

Chapter 5

Monitoring the System

This chapter describes the tasks you can do to check on the performance of the Barracuda Link Balancer. This section covers the following topics:

<i>Checking Status</i>	47
<i>Viewing Logs</i>	47
<i>Using a Syslog Server to Centrally Monitor System Logs</i>	48
<i>SNMP Monitoring</i>	48
<i>System Reports</i>	49
<i>Viewing System Tasks</i>	50

Checking Status

Check the **Basic > Status** page for an overview of the health and performance of your Barracuda Link Balancer, including:

- Utilization and status of the links.
- The subscription status of Energize Updates.
- System and hardware statistics, including CPU temperature and system load. Performance statistics displayed in red signify that the value exceeds the normal threshold.
- Incoming and outgoing traffic statistics for each WAN link.

You can also view WAN link utilization and connection status by scrolling over the WAN port graphic on the **Basic > Links** page. View the status of VPN tunnels using the **Services > VPN** page.

Viewing Logs

The Barracuda Link Balancer provides three types of logs under the **Logs** tab:

- Event Log - general system events.
- Firewall Log - firewall events.
- VPN Log - information about VPN tunnels.

Using the Web user interface, you can delete the log, filter the log entries that are displayed or export them to a CSV file.

View the system log displayed on the **Logs > Event Log** page to see events that have occurred. These include:

- Link status - A WAN link has become active or gone down; a link could not be detected.
- DHCP events - An IP address was handed out.
- Failed login attempts.

If the Barracuda Link Balancer firewall is enabled, you can view the firewall log on the **Logs > Firewall Log** page to see rules that have been executed and whether the traffic was dropped or allowed. Only rules that have the **Log** check box selected in their rule entry (under the **Firewall** tab) are logged in this way.

Check recent VPN tunnel activity by using the **Logs > VPN Log** page.

When any of these logs reaches their predetermined size a new log is started.

To have these logs emailed or sent to an FTP or SMB server on a regular basis, use the **Basic > Reports** page (models 330 and above).

Using a Syslog Server to Centrally Monitor System Logs

Syslog is a standard UNIX/Linux tool for logging messages and is available on all UNIX/Linux systems. The Barracuda Link Balancer writes to the syslog for link and system events. Use the **Advanced > Syslog** page to specify servers to which syslog data is sent.

SNMP Monitoring

Your SNMP monitor or other network management program can query the Barracuda Link Balancer SNMP agent for WAN link traffic statistics, amount of traffic going to and from the LAN, and hardware status. You can also receive SNMP traps that are generated if the WAN links become unavailable or if the Barracuda Link Balancer exceeds certain thresholds such as disk space usage.

To allow SNMP access to the Barracuda Link Balancer, navigate to the **Basic > Administration** page. On that page you can:

- Configure the Barracuda Link Balancer to accept and respond to SNMP queries.
- Update the SNMP community string.
- Set the SNMP version. Version 2c and version 3 are supported.
- Enter a range of IP addresses that are allowed to connect to the Barracuda Link Balancer using SNMP.
- Configure IP addresses that will be sent SNMP traps.

An SNMP monitor can access the Barracuda Link Balancer via any of the WAN or LAN IP addresses, although using the LAN is recommended in case one of the WAN links goes down.

Obtain and import these two MIB files to your SNMP monitor:

- The Barracuda Link Balancer MIB
- The Barracuda Reference MIB (standard across all Barracuda Networks products).

The MIB files are located on the Barracuda Link Balancer and can be obtained by replacing [LB IP] in the following URLs with a management IP address of your Barracuda Link Balancer:

- [http://\[LB IP\]:8000/Barracuda-BWB-MIB.txt](http://[LB IP]:8000/Barracuda-BWB-MIB.txt)
- [http://\[LB IP\]:8000/Barracuda-REF-MIB.txt](http://[LB IP]:8000/Barracuda-REF-MIB.txt)

SNMP Traps

An SNMP trap is generated by the Barracuda Link Balancer SNMP agent every five minutes if one of the following conditions is noted:

- CPU temperature exceeded its threshold.
- System temperature exceeded its threshold.
- CPU fan is dead.
- System fan is dead.
- Firmware storage exceeded its threshold.
- Log storage utilization exceeded its threshold.
- WANx is down.
- WANx is up.
- WANx reached configured saturation threshold.
- A high availability state change occurred.

Traps are sent to the SNMP trap receivers that are specified on the **Basic > Administration** page.

When any of these events is first noted, an email alert is sent to the system alerts email address specified on the **Basic > Administration** page. If an error condition continues to be detected, an email is sent every hour to the same email address.

System Reports



Note: Reporting is only available on models 330 and above.

Use the **BASIC > Reports** page to choose from a variety of information that can help you keep track of activity performed by the Barracuda Link Balancer. You can either generate a report on-demand for instant viewing or you can automatically generate scheduled reports for later delivery.

Reports can include any of the following trends:

- The average bandwidth usage by hour
- The total traffic by date
- The total link uptime by date
- The average VPN bandwidth usage by hour
- The VPN traffic by date
- The average TCP connections per hour
- The TCP connections by date

Reports can also include any of the following logs:

- Inbound link balancing
- Firewall activity
- VPN activity
- Link failover events
- Device failover events

The **Report Options** section allows you to choose the criteria for compiling the report data, as well as layout and output options. You can define a time frame for the report, select the interface links and VPN tunnels to include, and choose to analyze inbound traffic, outbound traffic, or both. It is also possible to select a layout for the graphical charts (available are lines, horizontal or vertical bars, or pie charts) as well as one of the offered output formats HTML, PDF, plain text, or comma-separated CSV.

Trend graphs can be selected to be included into the report as well as a choice of activity log summaries.



Note: If any VPN information was selected to be included with the report, the report can not be scheduled or executed without adding at least one VPN tunnel.

After making your choices, you may either execute report generation at once by saving the screen now, or you may schedule it for later and/or repeating execution by filling in the fields in the **Schedule Report** section. There, you start by filling in a report group name, followed by selecting your delivery options where you may either choose e-mail as the transport method, or an external server for FTPing or SMBing it. If you choose the latter, you will see a couple more fields in which you must provide the external server's IP address or hostname and the user credentials.

Once a report was scheduled, it will be listed in the **Scheduled Reports** section below from where it can be edited, disabled or deleted.

Viewing System Tasks

Go to the **Advanced > Task Manager** page to see a list of tasks that are in the process of being performed and any errors encountered when performing these tasks. Background tasks include firmware download and configuration restoration.

Chapter 6

Maintaining the Barracuda Link Balancer

This chapter describes how to maintain the Barracuda Link Balancer. The following topics are covered:

<i>Backing up and Restoring Your System Configuration</i>	51
<i>Updating the Firmware of Your Barracuda Link Balancer</i>	51
<i>Replacing a Failed System</i>	51
<i>Reloading, Restarting, and Shutting Down the System</i>	52
<i>Using the Built-in Troubleshooting Tools</i>	53
<i>Rebooting the System in Recovery Mode</i>	53

Backing up and Restoring Your System Configuration

Back up and restore the configuration of your Barracuda Link Balancer using the **Advanced > Backup** page. You should back up your system on a regular basis in case you need to restore this information on a replacement Barracuda Link Balancer or in the event the current system data becomes corrupt.

If you are restoring a backup file on a new Barracuda Link Balancer that is not configured, first enter the new system's IP address and DNS information on the **Basic > IP Configuration** page.

The following information is not included in the backup file:

- System password
- System IP information
- DNS information

Updating the Firmware of Your Barracuda Link Balancer

The **Advanced > Firmware Update** page allows you to manually update the firmware version of the system or revert to a previous version. The only time you should revert back to an old firmware version is if you recently downloaded a new version that is causing unexpected problems. In this case, call Barracuda Networks Technical Support before reverting back to a previous firmware version.

If there is a more recent firmware version available than what is already installed, the **Download Now** button will be enabled.

Applying a new firmware version will result in a short service outage.

Replacing a Failed System

Before you replace your Barracuda Link Balancer, use the tools provided on the **Advanced > Troubleshooting** page to try to resolve the problem.

In the event that a Barracuda Link Balancer fails and you cannot resolve the issue, customers that have purchased the Instant Replacement service can call Technical Support and arrange for a new unit to be shipped out within 24 hours.

After receiving the new system, ship the old Barracuda Link Balancer back to Barracuda Networks at the address below with an RMA number marked clearly on the package. Barracuda Networks Technical Support can provide details on the best way to return the unit.

Barracuda Networks
3175 S. Winchester Blvd.
Campbell, CA 95008



Note: To set up the new Barracuda Link Balancer so it has the same configuration as your old failed system, restore the backup file from the old system onto the new system, and then manually configure the new system's IP information on the **Basic > IP Configuration** page. For information on restoring data, refer to *Backing up and Restoring Your System Configuration* on page 51.

Reloading, Restarting, and Shutting Down the System

The **System Reload/Shutdown** section on the **Basic > Administration** page allows you to shutdown, restart, and reload system configuration on the Barracuda Link Balancer.

Shutting down the system powers off the unit. Restarting the system reboots the unit. Reloading the system re-applies the system configuration.

You can also reboot the Barracuda Link Balancer by pressing **RESET** on the front panel of the Barracuda Link Balancer.

Do not press and hold the **RESET** button for more than a couple of seconds. Holding it for five seconds or longer changes the IP address of the system. See *Using the Reset Button to Reset the LAN IP address* on page 52 for more information.

Using the Reset Button to Reset the LAN IP address

The Barracuda Link Balancer is assigned a default LAN IP address of 192.168.200.200. You can change this IP address in one of three ways:

- navigate using the Web user interface to the **Basic > IP Configuration** page
- connect a VGA monitor and a keyboard to the back of the Barracuda Link Balancer and using the Administrative Console (username **admin**, password **admin**)
- or press the **RESET** button on the front panel.

Pressing **RESET** for five seconds sets the LAN IP address to 192.168.200.200. Pressing **RESET** eight seconds changes the LAN IP address to 192.168.1.200. Pressing the button for 12 seconds changes the LAN IP address to 10.1.1.200. You will notice the three LEDs on the front panel flash at the same time intervals.

Using the Built-in Troubleshooting Tools

The **Advanced > Troubleshooting** page provides various tools that help troubleshoot network connectivity issues that may be impacting the performance of your Barracuda Link Balancer. You can perform a number of connectivity tests such as ping, telnet, dig/nslookup, TCP dump, and traceroute.

Barracuda Networks Technical Support may ask you to make a connection to Barracuda Central so they can help diagnose problems on your system.

Rebooting the System in Recovery Mode

If your Barracuda Link Balancer experiences a serious issue that impacts its core functionality, you can use diagnostic and recovery tools that are available at the reboot menu to return your system to an operational state.

Before you use the diagnostic and recovery tools, do the following:

- Use the built-in troubleshooting tools on the **Advanced > Troubleshooting** page to help diagnose the problem.
- Perform a system restore from the last known good backup file.
- Contact Barracuda Networks Technical Support for additional troubleshooting tips.

As a last resort, you can reboot your Barracuda Link Balancer and run a memory test or perform a complete system recovery, as described in this section.

To perform a system recovery or hardware test:

1. Connect a monitor and keyboard directly to your Barracuda Link Balancer.
2. Reboot the system by doing one of the following:
 - Click **Restart** on the **Basic > Administration** page.
 - Press the Power button on the front panel to turn off the system, and then press the Power button again to turn the system back on.

The Barracuda splash screen displays with the following three boot options:

Barracuda
Recovery
Hardware_Test

3. Use your keyboard to select the desired boot option, and click **Enter**.

You must select the boot option within three seconds of the splash screen appearing. If you do not select an option within three seconds, the Barracuda Link Balancer defaults to starting up in the normal mode (first option).

For a description of each boot option, refer to *Reboot Options* on page 54.

Reboot Options

Table 6.1 describes the options available at the reboot menu.

Table 6.1: Reboot Options

Reboot Options	Description
Barracuda	Starts the Barracuda Link Balancer in the normal (default) mode. This option is automatically selected if no other option is specified within the first three (3) seconds of the splash screen appearing.
Recovery	Displays the Recovery Console where you can select the following options: <ul style="list-style-type: none">• Perform file system repair—Repairs the file system on the Barracuda Link Balancer.• Perform full system re-image—Restores the factory settings on your Barracuda Link Balancer and clears out all configuration information.• Enable remote administration—Initiates a connection to Barracuda Central that allows Barracuda Networks Technical Support to access the system. Another method for enabling this troubleshooting connection is to click Establish Connection to Barracuda Central on the Advanced > Troubleshooting page.• Run diagnostic memory test—Runs a diagnostic memory test from the operating system. If problems are reported when running this option, we recommend running the Hardware_Test option next.
Hardware_Test	Performs a thorough memory test that shows most memory related errors within a two-hour time period. The memory test is performed outside of the operating system and can take a long time to complete. Reboot your Barracuda Link Balancer to stop the hardware test. You may do this by pressing Ctrl-Alt-Del on the keyboard, or by pressing the RESET button on the Barracuda Link Balancer.

Limited Warranty and License

Barracuda Networks Limited Hardware Warranty (v 2.1)

Barracuda Networks, Inc., or the Barracuda Networks, Inc. subsidiary or authorized Distributor selling the Barracuda Networks product, if sale is not directly by Barracuda Networks, Inc., ("Barracuda Networks") warrants that commencing from the date of delivery to Customer (but in case of resale by a Barracuda Networks reseller, commencing not more than sixty (60) days after original shipment by Barracuda Networks, Inc.), and continuing for a period of one (1) year: (a) its products (excluding any software) will be free from material defects in materials and workmanship under normal use; and (b) the software provided in connection with its products, including any software contained or embedded in such products will substantially conform to Barracuda Networks published specifications in effect as of the date of manufacture. Except for the foregoing, the software is provided as is. In no event does Barracuda Networks warrant that the software is error free or that Customer will be able to operate the software without problems or interruptions. In addition, due to the continual development of new techniques for intruding upon and attacking networks, Barracuda Networks does not warrant that the software or any equipment, system or network on which the software is used will be free of vulnerability to intrusion or attack. The limited warranty extends only to you the original buyer of the Barracuda Networks product and is non-transferable.

Exclusive Remedy

Your sole and exclusive remedy and the entire liability of Barracuda Networks under this limited warranty shall be, at Barracuda Networks or its service centers option and expense, the repair, replacement or refund of the purchase price of any products sold which do not comply with this warranty. Hardware replaced under the terms of this limited warranty may be refurbished or new equipment substituted at Barracuda Networks' option. Barracuda Networks obligations hereunder are conditioned upon the return of affected articles in accordance with Barracuda Networks then-current Return Material Authorization ("RMA") procedures. All parts will be new or refurbished, at Barracuda Networks' discretion, and shall be furnished on an exchange basis. All parts removed for replacement will become the property of Barracuda Networks. In connection with warranty services hereunder, Barracuda Networks may at its discretion modify the hardware of the product at no cost to you to improve its reliability or performance. The warranty period is not extended if Barracuda Networks repairs or replaces a warranted product or any parts. Barracuda Networks may change the availability of limited warranties, at its discretion, but any changes will not be retroactive. IN NO EVENT SHALL BARRACUDA NETWORKS LIABILITY EXCEED THE PRICE PAID FOR THE PRODUCT FROM DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES RESULTING FROM THE USE OF THE PRODUCT, ITS ACCOMPANYING SOFTWARE, OR ITS DOCUMENTATION.

Exclusions and Restrictions

This limited warranty does not apply to Barracuda Networks products that are or have been (a) marked or identified as "sample" or "beta," (b) loaned or provided to you at no cost, (c) sold "as is," (d) repaired, altered or modified except by Barracuda Networks, (e) not installed, operated or

maintained in accordance with instructions supplied by Barracuda Networks, or (f) subjected to abnormal physical or electrical stress, misuse, negligence or to an accident.

EXCEPT FOR THE ABOVE WARRANTY, BARRACUDA NETWORKS MAKES NO OTHER WARRANTY, EXPRESS, IMPLIED OR STATUTORY, WITH RESPECT TO BARRACUDA NETWORKS PRODUCTS, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTY OF TITLE, AVAILABILITY, RELIABILITY, USEFULNESS, MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NONINFRINGEMENT, OR ARISING FROM COURSE OF PERFORMANCE, DEALING, USAGE OR TRADE. EXCEPT FOR THE ABOVE WARRANTY, BARRACUDA NETWORKS' PRODUCTS AND THE SOFTWARE ARE PROVIDED "AS-IS" AND BARRACUDA NETWORKS DOES NOT WARRANT THAT ITS PRODUCTS WILL MEET YOUR REQUIREMENTS OR BE UNINTERRUPTED, TIMELY, AVAILABLE, SECURE OR ERROR FREE, OR THAT ANY ERRORS IN ITS PRODUCTS OR THE SOFTWARE WILL BE CORRECTED. FURTHERMORE, BARRACUDA NETWORKS DOES NOT WARRANT THAT BARRACUDA NETWORKS PRODUCTS, THE SOFTWARE OR ANY EQUIPMENT, SYSTEM OR NETWORK ON WHICH BARRACUDA NETWORKS PRODUCTS WILL BE USED WILL BE FREE OF VULNERABILITY TO INTRUSION OR ATTACK.

Barracuda Networks Software License Agreement (v 2.1)

PLEASE READ THIS SOFTWARE LICENSE AGREEMENT ("AGREEMENT") CAREFULLY BEFORE USING THE BARRACUDA NETWORKS SOFTWARE. BY USING THE BARRACUDA SOFTWARE YOU ARE AGREEING TO BE BOUND BY THE TERMS OF THIS LICENSE. IF YOU ARE A CORPORATION, PARTNERSHIP OR SIMILAR ENTITY, THEN THE SOFTWARE LICENSE GRANTED UNDER THIS AGREEMENT IS EXPRESSLY CONDITIONED UPON ACCEPTANCE BY A PERSON WHO IS AUTHORIZED TO SIGN FOR AND BIND THE ENTITY. IF YOU ARE NOT AUTHORIZED TO SIGN FOR AND BIND THE ENTITY OR DO NOT AGREE WITH ALL THE TERMS OF THIS AGREEMENT, DO NOT USE THE SOFTWARE. IF YOU DO NOT AGREE TO THE TERMS OF THIS LICENSE YOU MAY RETURN THE SOFTWARE OR HARDWARE CONTAINING THE SOFTWARE FOR A FULL REFUND TO YOUR PLACE OF PURCHASE.

1. The software and documentation, whether on disk, in flash memory, in read only memory, or on any other media or in any other form (collectively "Barracuda Software") is licensed, not sold, to you by Barracuda Networks, Inc. ("Barracuda") for use only under the terms of this Agreement, and Barracuda reserves all rights not expressly granted to you. The rights granted are limited to Barracuda's intellectual property rights in the Barracuda Software and do not include any other patent or intellectual property rights. You own the media on which the Software is recorded but Barracuda retains ownership of the Software itself. If you have not completed a purchase of the Software and made payment for the purchase, the Software may only be used for evaluation purposes and may not be used in any production capacity. Furthermore the Software, when used for evaluation, may not be secure and may use publically available passwords.
2. Permitted License Uses and Restrictions. If you have purchased a Barracuda Networks hardware product, this Agreement allows you to use the Software only on the single Barracuda labeled hardware device on which the software was delivered. You may not make copies of the Software. You may not make a backup copy of the Software. If you have purchased a Barracuda Networks Virtual Machine you may use the software only in the licensed number of instances of the licensed sizes and you may not exceed the licensed capacities. You may make a reasonable number of backup copies of the Software. If you have purchased client software you may install the software only on the number of licensed clients. You may make a reasonable number of backup copies of the Software. For all purchases you may not modify or create derivative works of the Software except as provided by the Open Source Licenses included below. You may not make the Software available over a

network where it could be utilized by multiple devices or copied. Unless otherwise expressly provided in the documentation, your use of the Software shall be limited to use on a single hardware chassis, on a single central processing unit, as applicable, or use on such greater number of chassis or central processing units as you may have paid Barracuda Networks the required license fee; and your use of the Software shall also be limited, as applicable and set forth in your purchase order or in Barracuda Networks' product catalog, user documentation, or web site, to a maximum number of (a) seats (i.e. users with access to install Software), (b) concurrent users, sessions, ports, and/or issued and outstanding IP addresses, and/or (c) central processing unit cycles or instructions per second. Your use of the Software shall also be limited by any other restrictions set forth in your purchase order or in Barracuda Networks' product catalog, user documentation or Web site for the Software. The BARRACUDA SOFTWARE IS NOT INTENDED FOR USE IN THE OPERATION OF NUCLEAR FACILITIES, AIRCRAFT NAVIGATION OR COMMUNICATION SYSTEMS, LIFE SUPPORT MACHINES, OR OTHER EQUIPMENT IN WHICH FAILURE COULD LEAD TO DEATH, PERSONAL INJURY, OR ENVIRONMENTAL DAMAGE. YOU EXPRESSLY AGREE NOT TO USE IT IN ANY OF THESE OPERATIONS.

3. You may not transfer, rent, lease, lend, or sublicense the Software or allow a third party to do so. YOU MAY NOT OTHERWISE TRANSFER THE SOFTWARE OR ANY OF YOUR RIGHTS AND OBLIGATIONS UNDER THIS AGREEMENT. You agree that you will have no right and will not, nor will it assist others to: (i) make unauthorized copies of all or any portion of the Software; (ii) sell, sublicense, distribute, rent or lease the Software; (iii) use the Software on a service bureau, time sharing basis or other remote access system whereby third parties other than you can use or benefit from the use of the Software; (iv) disassemble, reverse engineer, modify, translate, alter, decompile or otherwise attempt to discern the source code of all or any portion of the Software; (v) utilize or run the Software on more computers than you have purchased license to; (vi) operate the Software in a fashion that exceeds the capacity or capabilities that were purchased by you.

4. THIS AGREEMENT SHALL BE EFFECTIVE UPON INSTALLATION OF THE SOFTWARE OR PRODUCT AND SHALL TERMINATE UPON THE EARLIER OF: (A) YOUR FAILURE TO COMPLY WITH ANY TERM OF THIS AGREEMENT OR (B) RETURN, DESTRUCTION OR DELETION OF ALL COPIES OF THE SOFTWARE IN YOUR POSSESSION. Rights of Barracuda Networks and your obligations shall survive any termination of this Agreement. Upon termination of this Agreement by Barracuda Networks, You shall certify in writing to Barracuda Networks that all copies of the Software have been destroyed or deleted from any of your computer libraries, storage devices, or any other location.

5. YOU EXPRESSLY ACKNOWLEDGE AND AGREE THAT THE USE OF THE BARRACUDA SOFTWARE IS AT YOUR OWN RISK AND THAT THE ENTIRE RISK AS TO SATISFACTION, QUALITY, PERFORMANCE, AND ACCURACY IS WITH YOU. THE BARRACUDA SOFTWARE IS PROVIDED "AS IS" WITH ALL FAULTS AND WITHOUT WARRANTY OF ANY KIND, AND BARRACUDA HEREBY DISCLAIMS ALL WARRANTIES AND CONDITIONS WITH RESPECT TO THE BARRACUDA SOFTWARE, EITHER EXPRESSED OR IMPLIED OR STATUTORY, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES AND/OR CONDITIONS OF MERCHANTABILITY, OF SATISFACTORY QUALITY, OF FITNESS FOR ANY APPLICATION, OF ACCURACY, AND OF NON-INFRINGEMENT OF THIRD PARTY RIGHTS. BARRACUDA DOES NOT WARRANT THE CONTINUED OPERATION OF THE SOFTWARE, THAT THE PERFORMANCE WILL MEET YOUR EXPECTATIONS, THAT THE FUNCTIONS WILL MEET YOUR REQUIREMENTS, THAT THE OPERATION WILL BE ERROR FREE OR CONTINUOUS, THAT CURRENT OR FUTURE VERSIONS OF ANY OPERATING SYSTEM WILL BE SUPPORTED, OR THAT DEFECTS WILL BE CORRECTED. NO ORAL OR WRITTEN INFORMATION GIVEN BY BARRACUDA OR AUTHORIZED BARRACUDA REPRESENTATIVE SHALL CREATE A WARRANTY. SHOULD THE BARRACUDA SOFTWARE PROVE DEFECTIVE, YOU ASSUME THE ENTIRE COST OF ALL NECESSARY SERVICING, REPAIR, OR CORRECTION. FURTHERMORE BARRACUDA NETWORKS SHALL ASSUME NO WARRANTY FOR ERRORS/BUGS, FAILURES OR DAMAGE WHICH

WERE CAUSED BY IMPROPER OPERATION, USE OF UNSUITABLE RESOURCES, ABNORMAL OPERATING CONDITIONS (IN PARTICULAR DEVIATIONS FROM THE INSTALLATION CONDITIONS) AS WELL AS BY TRANSPORTATION DAMAGE. IN ADDITION, DUE TO THE CONTINUAL DEVELOPMENT OF NEW TECHNIQUES FOR INTRUDING UPON AND ATTACKING NETWORKS, BARRACUDA NETWORKS DOES NOT WARRANT THAT THE SOFTWARE OR ANY EQUIPMENT, SYSTEM OR NETWORK ON WHICH THE SOFTWARE IS USED WILL BE FREE OF VULNERABILITY TO INTRUSION OR ATTACK. YOU EXPRESSLY ACKNOWLEDGE AND AGREE THAT YOU WILL PROVIDE AN UNLIMITED PERPETUAL ZERO COST LICENSE TO BARRACUDA FOR ANY PATENTS OR OTHER INTELLECTUAL PROPERTY RIGHTS WHICH YOU EITHER OWN OR CONTROL THAT ARE UTILIZED IN ANY BARRACUDA PRODUCT.

6. Termination and Fair Use Policy. BARRACUDA SHALL HAVE THE ABSOLUTE AND UNILATERAL RIGHT AT ITS SOLE DISCRETION TO DENY USE OF, OR ACCESS TO BARRACUDA SOFTWARE, IF YOU ARE DEEMED BY BARRACUDA TO BE USING THE SOFTWARE IN A MANNER NOT REASONABLY INTENDED BY BARRACUDA OR IN VIOLATION OF ANY LAW.

7. Limitation of Liability. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT SHALL BARRACUDA BE LIABLE FOR PERSONAL INJURY OR ANY INCIDENTAL, SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES WHATSOEVER, INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, LOSS OF DATA, BUSINESS INTERRUPTION, OR ANY OTHER COMMERCIAL DAMAGES OR LOSSES, ARISING OUT OF OR RELATED TO YOUR ABILITY TO USE OR INABILITY TO USE THE BARRACUDA SOFTWARE HOWEVER CAUSED, REGARDLESS OF THE THEORY OF LIABILITY AND EVEN IF BARRACUDA HAS BEEN ADVISED OF THE POSSIBILITY OF DAMAGES. In no event shall Barracuda's total liability to you for all damages exceed the amount of one hundred dollars. The following terms govern your use of the Energize Update Software except to the extent a particular program (a) is the subject of a separate written agreement with Barracuda Networks or (b) includes a separate "click-on" license agreement as part of the installation and/or download process. To the extent of a conflict between the provisions of the foregoing documents, the order of precedence shall be (1) the written agreement, (2) the click-on agreement, and (3) this Energize Update Software License.

8. Content Restrictions. YOU MAY NOT (AND MAY NOT ALLOW A THIRD PARTY TO) COPY, REPRODUCE, CAPTURE, STORE, RETRANSMIT, DISTRIBUTE, OR BURN TO CD (OR ANY OTHER MEDIUM) ANY COPYRIGHTED CONTENT THAT YOU ACCESS OR RECEIVE THROUGH USE OF THE PRODUCT CONTAINING THE SOFTWARE. YOU ASSUME ALL RISK AND LIABILITY FOR ANY SUCH PROHIBITED USE OF COPYRIGHTED CONTENT. You agree not to publish any benchmarks, measurements, or reports on the product without Barracuda Networks' written express approval.

9. Third Party Software. Some Software which supports Bare Metal Disaster Recovery of Microsoft Windows Vista and Microsoft Windows 2008 Operating Systems (DR6) contains and uses components of the Microsoft Windows Pre-Installation Environment (WINPE) with the following restrictions: (i) the WINPE components in the DR6 product are licensed and not sold and may only be used with the DR6 product; (ii) DR6 is provided "as is"; (iii) Barracuda and its suppliers reserve all rights not expressly granted; (iv) license to use DR6 and the WINPE components is limited to use of the product as a recovery utility program only and not for use as a general purpose operating system; (v) Reverse engineering, decompiling or disassembly of the WINPE components, except to the extent expressly permitted by applicable law, is prohibited; (vi) DR6 contains a security feature from Microsoft that will automatically reboot the system without warning after 24 hours of continuous use; (vii) Barracuda alone will provide support for customer issues with DR6 and Microsoft and its Affiliates are released of all liability related to its use and operation; and, (viii) DR6 is subject to U.S. export jurisdiction.

10. Trademarks. Certain portions of the product and names used in this Agreement, the Software and the documentation may constitute trademarks of Barracuda Networks. You are not authorized to use any such trademarks for any purpose.

11. Export Restrictions. You may not export or re-export the Software without: (a) the prior written consent of Barracuda Networks, (b) complying with applicable export control laws, including, but not limited to, restrictions and regulations of the Department of Commerce or other United States agency or authority and the applicable EU directives, and (c) obtaining any necessary permits and licenses. In any event, you may not transfer or authorize the transfer of the Software to a prohibited territory or country or otherwise in violation of any applicable restrictions or regulations. If you are a United States Government agency the Software and documentation qualify as "commercial items", as that term is defined at Federal Acquisition Regulation ("FAR") (48 C.F.R.) 2.101, consisting of "commercial computer software" and "commercial computer software documentation" as such terms are used in FAR 12.212. Consistent with FAR 12.212 and DoD FAR Supp. 227.7202-1 through 227.7202-4, and notwithstanding any other FAR or other contractual clause to the contrary in any agreement into which this Agreement may be incorporated, Government end user will acquire the Software and documentation with only those rights set forth in this Agreement. Use of either the Software or documentation or both constitutes agreement by the Government that the Software and documentation are "commercial computer software" and "commercial computer software documentation", and constitutes acceptance of the rights and restrictions herein.

12. General. THIS AGREEMENT IS GOVERNED BY THE LAWS OF THE STATE OF CALIFORNIA, USA WITH JURISDICTION OF SANTA CLARA COUNTY, CALIFORNIA, UNLESS YOUR HEADQUARTERS IS LOCATED IN SWITZERLAND, THE EU, OR JAPAN. IF YOUR HEADQUARTERS IS LOCATED IN SWITZERLAND THE SWISS MATERIAL LAW SHALL BE USED AND THE JURISDICTION SHALL BE ZURICH. IF YOUR HEADQUARTERS IS LOCATED IN THE EU, AUSTRIAN LAW SHALL BE USED AND JURISDICTION SHALL BE INNSBRUCK. IF YOUR HEADQUARTERS IS LOCATED IN JAPAN, JAPANESE LAW SHALL BE USED AND JURISDICTION SHALL BE TOKYO. THIS AGREEMENT WILL NOT BE SUBJECT TO ANY CONFLICT-OF-LAWS PRINCIPLES IN ANY JURISDICTION. THIS AGREEMENT WILL NOT BE GOVERNED BY THE U.N. CONVENTION ON CONTRACTS FOR THE INTERNATIONAL SALES OF GOODS. This Agreement is the entire agreement between You and Barracuda Networks regarding the subject matter herein and supersedes any other communications with respect to the Software. If any provision of this Agreement is held invalid or unenforceable, the remainder of this Agreement will continue in full force and effect. Failure to prosecute a party's rights with respect to a default hereunder will not constitute a waiver of the right to enforce rights with respect to the same or any other breach.

13. Assignability. You may not assign any rights or obligations hereunder without prior written consent from Barracuda Networks.

14. Billing Issues. You must notify Barracuda of any billing problems or discrepancies within sixty (60) days after they first appear on the statement you receive from your bank, Credit Card Company, other billing company or Barracuda Networks. If you do not bring such problems or discrepancies to Barracuda Networks attention within the sixty (60) day period, you agree that you waive the right to dispute such problems or discrepancies.

15. Collection of Data. You agree to allow Barracuda Networks to collect information ("Statistics") from the Software in order to fight spam, virus, and other threats as well as optimize and monitor the Software. Information will be collected electronically and automatically. Statistics include, but are not limited to, the number of messages processed, the number of messages that are categorized as spam, the number of virus and types, IP addresses of the largest spam senders, the number of emails classified for Bayesian analysis, capacity and usage, and other statistics. Your data will be kept private and will only be reported in aggregate by Barracuda Networks.

16. Subscriptions. Software updates and subscription information provided by Barracuda Energize Updates or other services may be necessary for the continued operation of the Software. You acknowledge that such a subscription may be necessary. Furthermore some functionality may only be available with additional subscription purchases. Obtaining Software updates on systems where no valid subscription has been purchased or obtaining functionality where subscription has not been purchased is strictly forbidden and in violation of this Agreement. All initial subscriptions commence at the time of activation and all renewals commence at the expiration of the previous valid subscription. Unless otherwise expressly provided in the documentation, you shall use the Energize Updates Service and other subscriptions solely as embedded in, for execution on, or (where the applicable documentation permits installation on non-Barracuda Networks equipment) for communication with Barracuda Networks equipment owned or leased by you. All subscriptions are non-transferrable. Barracuda Networks makes no warranty that subscriptions will continue uninterrupted. Subscription may be terminated without notice by Barracuda Networks for lack of full payment.
17. Auto Renewals. If your Software purchase is a time based license, includes software maintenance, or includes a subscription, you hereby agree to automatically renew this purchase when it expires unless you notify Barracuda 15 days before the renewal date. Barracuda Networks will automatically bill you or charge you unless notified 15 days before the renewal date.
18. Time Base License. If your Software purchase is a time based license you expressly acknowledge that the Software will stop functioning at the time the license expires. You expressly indemnify and hold harmless Barracuda Networks for any and all damages that may occur because of this.
19. Support. Telephone, email and other forms of support will be provided to you if you have purchased a product that includes support. The hours of support vary based on country and the type of support purchased. Barracuda Networks Energize Updates typically include Basic support.
20. Changes. Barracuda Networks reserves the right at any time not to release or to discontinue release of any Software or Subscription and to alter prices, features, specifications, capabilities, functions, licensing terms, release dates, general availability or other characteristics of any future releases of the Software or Subscriptions.
21. Open Source Licensing. Barracuda Networks products may include programs that are covered by the GNU General Public License (GPL) or other Open Source license agreements, in particular the Linux operating system. It is expressly put on record that the Software does not constitute an edited version or further development of the operating system. These programs are copyrighted by their authors or other parties, and the authors and copyright holders disclaim any warranty for such programs. Other programs are copyright by Barracuda Networks. Further details may be provided in an appendix to this agreement where the licenses are re-printed. Barracuda Networks makes available the source code used to build Barracuda products available at source.barracuda.com. This directory includes all the programs that are distributed on the Barracuda products. Obviously not all of these programs are utilized, but since they are distributed on the Barracuda product we are required to make the source code available.

Barracuda Networks Energize Updates and Other Subscription Terms

Barracuda Networks Software License Agreement Appendix

The GNU General Public License (GPL) Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program

is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement).

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF GNU TERMS AND CONDITIONS

Barracuda Networks Products may contain programs that are copyright (c)1995-2005 International Business Machines Corporation and others. All rights reserved. These programs are covered by the following License: "Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice appear in all copies of the

Software and that both the above copyright notice(s) and this permission notice appear in supporting documentation."

Barracuda Networks Products may include programs that are covered by the BSD License: "Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

The names of the authors may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE."

Barracuda Networks Products may include the libspf library which is Copyright (c) 2004 James Couzens & Sean Comeau, All rights reserved. It is covered by the following agreement: Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: 1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. 2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. **THIS SOFTWARE IS PROVIDED "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED.** IN NO EVENT SHALL THE AUTHORS MAKING USE OF THIS LICENSE OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Barracuda Networks Products may contain programs that are Copyright (c) 1998-2003 Carnegie Mellon University. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: 1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. 2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. The name "Carnegie Mellon University" must not be used to endorse or promote products derived from this software without prior written permission. For permission or any other legal details, please contact Office of Technology Transfer, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213-3890 (412) 268-4387, fax: (412) 268-7395, tech-transfer@andrew.cmu.edu. Redistributions of any form whatsoever must retain the following acknowledgment: "This product includes software developed by Computing Services at Carnegie Mellon University (<http://www.cmu.edu/computing/>)."**CARNEGIE MELLON UNIVERSITY DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS, AND IN NO EVENT SHALL CARNEGIE MELLON UNIVERSITY BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT,**

NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Barracuda Networks Software may include programs that are covered by the Apache License or other Open Source license agreements. The Apache license is re-printed below for you reference. These programs are copyrighted by their authors or other parties, and the authors and copyright holders disclaim any warranty for such programs. Other programs are copyright by Barracuda Networks.

Apache License

Version 2.0, January 2004

<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the

purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

- (a) You must give any other recipients of the Work or Derivative Works a copy of this License; and
- (b) You must cause any modified files to carry prominent notices stating that You changed the files; and
- (c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
- (d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.

6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

Barracuda Networks makes available the source code used to build Barracuda products available at source.barracuda.com. This directory includes all the programs that are distributed on the Barracuda products. Obviously not all of these programs are utilized, but since they are distributed on the Barracuda product we are required to make the source code available.

Index

Numerics

1:1 NAT [4](#), [22](#)

A

additional IP addresses

- in front of your firewall [13](#)
- replacing your firewall [17](#)

Administration page [36](#), [48](#), [52](#)

alerts [36](#)

applications

- creating custom [24](#)
- definition [24](#)

ARIN networks [26](#), [33](#)

B

backing up configuration [51](#)

backup link [24](#)

Backup page [51](#)

C

character tags [55](#)

checking status [47](#)

configuring network settings [19](#)

contacting technical support [6](#)

D

default IP address [52](#)

denial of service attacks [22](#)

DHCP server [21](#)

diagnostic memory test [54](#)

DNS [21](#)

DNS servers, configuring [21](#)

E

externally accessible IP addresses [26](#)

F

failed system, replacing [52](#)

H

hardware test [54](#)

I

IP aliases [21](#)

IP masquerading [4](#)

L

LAN IP address, in front of your firewall [9](#), [13](#)

M

Many to 1 NAT [4](#), [22](#)

masquerade IP address [27](#)

N

NAT/Port Forwarding [20](#), [22](#)

NAT-T, NAT-Traversal [28](#)

notifications [36](#)

O

online help [2](#)

P

persistence [4](#)

Port Address Translation [4](#), [22](#)

Port forwarding [22](#)

port forwarding [4](#)

port forwarding rules [23](#)

primary link [24](#)

private link [24](#)

Q

Quality of Service rules, configuring [25](#)

R

reboot options [53](#)

recovery mode [53](#)

re-imaging system [54](#)

reloading the system [52](#)

remote administration [54](#)

repairing, file system [54](#)

replacing failed system [52](#)

reporting, system report, statistics [49](#)

RESET button, using [52](#)

restarting the system [52](#)

restoring configuration [51](#)

S

shutting down the system [52](#)

SNMP [48](#), [49](#)

community string [48](#)

MIB [48](#)

SNMP traps

configure trap receivers [48](#)

statistics, system report [49](#)

Status page [47](#)
system alerts email address [49](#)

T

Task Manager page [50](#)
technical support, contacting [6](#)
testing memory [54](#)
time zone, setting [36](#)
Troubleshooting page [53](#)

U

updating firmware [51](#)

V

virtual interfaces [21](#)
VPN tunnel
 failover [29](#)
 shared secret [28](#)
 SSL certificates [28](#)
VPN Tunnel as Failover Link for a Broken MPLS Link [29](#)
VPN Tunnel, failover link for broken site-to-site WAN link
 [29](#)
VPN Tunnel, troubleshooting [29](#)

W

WAN interface, Web user interface access via [36](#)
WAN IP impersonation [20](#)
Web user interface
 logging in [12, 16](#)
weight, WAN link [24](#)